

hp·ux/usr

Hands-On Solutions for HP-UX Users • November/December 1996



Protect Your Computer As You Protect Your Business
The SpoolKick Procedure ■ Sockets Programming

A Publication of Interex • The International Association of Hewlett-Packard Computing Professionals

BULK RATE
US Postage
PAID
Saratoga, MS
Permit #30

GALAXY 1100

PORTABLE

HIGH-PERFORMANCE UNIX WORKSTATION



The Galaxy 1100

Complex applications demand sophisticated tools. Whether it be software development and demonstration, network administration, or other mobile HP-UX computing requirements, SAIC's Galaxy 1100 portable workstation does the job. Galaxy delivers critical information to the commercial, industrial, and military user with the performance and reliability of the HP PA-RISC processor.

Fully software compatible with the HP 9000/700 workstation series, this powerful and portable workstation offers superb graphics capability. Its 10.4-inch, full-color active matrix LCD renders images at a resolution of 1024 x 768. With minimum processing power of 79 SPECfp92 and 122 SPECint92, the Galaxy is equipped to handle even the most demanding computations.

Key Features

- 80 MHz HP 712 PA-RISC CPU
- 32 MB of RAM expandable to 128 MB with error detection/correction capability
- 2 GB hard drive
- 3.5-inch 1.44 MB floppy drive
- 10.4-inch color active matrix LCD
- 2 integral PCMCIA slots
- IEEE 802.3 Ethernet interface
- SCSI-2 interface
- POSIX-based HP-UX operating system



Science Applications International Corporation
10770 Wateridge Circle, San Diego, CA 92121
Tel 619-552-5334 or 1-800-447-4373 Fax 619-552-5253
<http://www.saic.com/it/stp/galaxy/index.html>

CIRCLE 138 ON READER SERVICE CARD

Is It New Or Is It Refurbished



**The only difference
is the money you save.**

With 17 years of knowledgeable experience and reliable service, Monterey Bay Communications is a leader in Hewlett-Packard workstation remarketing. We're professionals at providing HP 1000 and 9000 users with reliable equipment that is functionally and cosmetically identical to what HP offers — and at substantial cost savings. In addition to the 700 / 400 / 300 / 200 series, Monterey Bay Communications also offers mass storage systems,



monitors, memory and interfaces, as well as a variety of printers.

All equipment is warranted and eligible for

Hewlett-Packard maintenance. An extensive parts and spares inventory and knowledgeable staff ensure prompt service and immediate delivery.

For more detailed information or a price quotation, give MBC a call at 408/429-6144.

	NEW	MBC
Performance	✓	✓
Warranty	✓	✓
Support / Tech Expertise	✓	✓
Accessories	✓	✓
Maintenance Eligibility	✓	✓
Substantial Cost Savings		✓
Simple Order Processing		✓
Immediate Delivery		✓

CIRCLE 7 ON READER SERVICE CARD

MONTEREY BAY

COMMUNICATIONS

The HP Workstation Remarketing Specialists

Monterey Bay Communications Inc., 1010 Fair Avenue, Santa Cruz, CA 95060 Tel: 408-429-6144 Fax: 408-429-1918

SOLID GOLD CONNECTIVITY



**Now you can access
any mainframe file
platform
application or
operating system.
From any desktop.**



PERICOM
US, UK, GERMANY, FRANCE

2271 State Highway 33, Suite 106
Hamilton Square, NJ 08690
Phone: (609) 588-5300
Fax: (609) 588-8906
BBS: (609) 895-0767
email: sales@pericom-usa.com
Sales: 1-800-233-2206

CIRCLE 97 ON READER SERVICE CARD

Data Movement Solved!

HP ORACLE UNIX

IMAGE MIGRATION

FLAT FILES CODASYL DBMS DEC VMS TEST DATABASES
ALLBASE ARCHIVING DECISION SUPPORT

Thanks, Warehouse.

Warehouse does two things. Move and manipulate data. Almost anywhere. No programming needed.

If you need to cross HP, UNIX or VAX platforms. Or move between Image, Oracle, DBMS and other databases or applications. No problem. If you need to migrate legacy data across the network straight into relational databases. Done. It arrives as native data.

Data warehousing is just as easy. In minutes, you can archive and retrieve selected data from any media. Make

test and decision support databases. Think of the advantages. No downtime. No big intermediate flat files. And total platform transparency.

Hundreds of companies depend on Warehouse. And Taurus is a certified "HP User Reference" company as well as a participant in the "Oracle Business Alliance Program." So if you move data, now there's a real solution. Warehouse from Taurus. Call today for details. Because Warehouse moves data.

Taurus Software


1032 Elwell Court, Palo Alto, CA 94303

Phone: 415-961-1323 x100 • Fax: 415-961-1454 • E-mail: sales@taurus.com • Web: www.taurus.com

CIRCLE 123 ON READER SERVICE CARD

Contents

Features



page 38

**Protect Your Computer as
You Protect Your Business**

by John A. Pezzano 22

The Spoolkick Procedure

by Bill Hassell 30

Sockets Programming

by Andrew J. Phillips 38

Book Review: The Dilbert Principle

by Chris Curtin 52

Departments

Q & A 8

CSL Perspective 64

HP 1000 Guru 66

New Products 74

Columns

HP-UX Systems Administration by Chris Curtin 16

HP-UX by David L. Totsch 20

Windows NT by Bob Combs 54

Internet Goodies by Joe Berry 56


X-Watch by Larry Headlund 60

HP-RT by Anita Harris 70

New Products
See page 74



Cover Story:
See page 22



Lose your mouse and increase your productivity.

From Wall Street to Silicon Valley, your top competitors have replaced hundreds of free mice that came with their workstations with \$199 MOUSE-TRAK™ trackballs. The reason: productivity and reliability.

Productivity: Laboratory testing has shown that only 4 hours of continuous mouse usage can result in as much as 60% loss of hand strength. A trader, engineer, or data entry user in that condition is simply not going to be as productive in the second half of the day as in the first. The same tests show no signs of physical fatigue when using a MOUSE-TRAK.


Reliability: MOUSE-TRAK's rugged construction results in *much* higher reliability than mice or consumer trackballs. MOUSE-TRAK doesn't take traders out of play or make support people pull their hair out!

Call, Fax, or email today to order or receive more information about MOUSE-TRAK.

1 - 8 0 0 - 5 3 3 - 4 8 2 2
sales @ moustrak.com
<http://www.mousetrak.com>

mouse-trak®

The Professional's Trackball



ITAC Systems, Inc. 214 494 3073 Fax 214 494-4159

International MOUSE-TRAK dealers: **UK** NTWare Ltd. Tel: 0865 784990 **SunExpress** Tel: 0800 89 88 88 • **France** PHASELYS Tel: 1 43 94 42 42 **RISC TECHNOLOGY EUROPE** Tel: 01 41 85 10 20 **SunExpress** Tel: 05 90 61 57 • **Germany** The Chameleon Group Tel: 0211-379057 **SunExpress** Tel: 01 30 81 61 91 • **Switzerland** Datacomp Tel: 01 740 51 40 **SunExpress** Tel: 155 19 26 • **Denmark** DeSeCAM Tel: 48 24 12 04 • **Holland** SunExpress Tel: 06 022 34 45 • **Canada** Guillevin Tel: 416 670 8433 • **Australia** Hypec Electronics Tel: (02) 808 3666 • **Sweden** SPECMA Tel: 031 89 16 00 **SunExpress** Tel: 020 795 726 • **Korea** Aurora Systems, Inc. Tel: 718-4985 • **Israel** Seg Tec Tel: 972 3 556 7458 • **Japan** SunExpress Tel: 0120-33-9096 MOUSE-TRAK is manufactured in the U.S.A. by ITAC Systems, Inc., 3113 Benton Street, Garland, Texas 75042 Fax: 214/494-4159 Patented by ITAC Systems, Inc. MOUSE-TRAK is a registered trademark of ITAC Systems, Inc. Other brand and product names are trademarks of their respective holders.

CIRCLE 117 ON READER SERVICE CARD

<http://www.interex.org/>

INTEREX EXECUTIVE DIRECTOR

Charles A. Piercey

PUBLICATIONS MANAGER

Connie Wright

EDITORIAL

MANAGING EDITOR

Michael Ehrhardt

COPY EDITOR

Richard Kranz

NEW PRODUCTS EDITOR

Michelle Pollace

EDITOR

Jean Nattkemper

ADVERTISING & MARKETING

ADVERTISING SALES MANAGER

Brian Hallin

ADVERTISING REPRESENTATIVE

Kathie Schwartz

ADVERTISING COORDINATOR

Phil Nguyen

DESIGN AND PRODUCTION

SENIOR GRAPHICS DESIGNER

Molly McGinnity

DESKTOP PUBLISHING SPECIALIST

June Ramirez

GRAPHICS PRODUCTION SPECIALIST

Gale Patterson

hp-ux/usr is published bimonthly by Interex, the International Association of Hewlett-Packard Computing Professionals. Second-class application pending at Sunnyvale, California 94086 and additional offices. The editorial and business offices are located at 1192 Borregas Ave., Sunnyvale, California 94089, USA, 408.747.0227, Fax 408.747.0947. Address membership questions and change of address to Membership Services. Address all questions concerning circulation/distribution to the Distribution Manager.

Remittances should be sent to Interex,
File No. 61054, P. O. Box 60000,
San Francisco, California 94160, USA.

Address all editorial correspondence to Michael Ehrhardt,
Editor, *hp-ux/usr* Magazine, c/o Interex,
P.O. Box 3439, Sunnyvale, California 94088-3439, USA.

Subscription to *hp-ux/usr* is \$49.50 (6 issues) per year in the U.S., add \$25.00 for Canada and Mexico, add \$50.00 for all other countries. Member Services (Associate, Contributing or Online Service Package membership levels) include a subscription to *hp-ux/usr* at \$49.50. For other Member Services refer to membership form.

Statements of fact and opinion are the responsibility of the authors alone and do not imply an opinion on the part of the Interex Board or Magazine. Entire contents copyright © 1996 by Interex. All rights reserved.

POSTMASTER: Send address changes to:
Interex, P.O. Box 3439,
Sunnyvale, California 94088-3439 USA.
Attention: Member Services.

TRADEMARKS: UNIX, X/Open Co. Ltd.;
HP-UX, Hewlett-Packard; X Window System,
X Consortium, Inc.

How to Contact Interex...

EDITORIAL SUBMISSIONS

hp-ux/usr encourages readers to contribute their opinions, tips, and solutions. When sending letters for publication or to request author contribution guidelines, please address them to *hp-ux/usr* editor Michael Ehrhardt.

Postal Address:

Interex
P.O. Box 3439
Sunnyvale, CA 94088-3439

Office Address:

1192 Borregas Avenue
Sunnyvale, CA 94089

Because of the difference in zip codes between our office address and P.O. Box, please be sure to address all regular mail to the P.O. Box. Any express service packages should be delivered to the Borregas Avenue address. Thank you for your attention to this small but significant detail.

TELEPHONE:

The Interex switchboard is open 8:00 a.m.–5:00 p.m., Pacific Time. Call 800.468.3739 (U.S. and Canada) or 408.747.0227. After 5:30 p.m. our voicemail system will record your call.

FACSIMILE:

Call 408.747.0947

INTERNET:

To send e-mail to Interex, use the following address format:

<IDname>@interex.org

The following IDs are currently active on Interex's HP 9000:

Address	Department
ehrhadt	<i>hp-ux/usr</i> Letters to the Editor, Q&A, and requests for author guidelines
webmaster	Internet support
cslhpx	Contributed Software Library
membership	Membership/subscription inquiries and services
pubs	Circulation and advertising inquiries
conference	Conference questions and arrangements

An example would be `pubs@interex.org`
Anything before the @ sign is case insensitive.

COMPUSERVE:

Interex can be contacted via the CompuServe ↔ Internet gateway. To send CompuServe mail, use the following format:

>INTERNET: `pubs@interex.org`

You can address your mail to specific departments using the ID's listed above.

Interex maintains a CompuServe account that is collected daily. Please address all messages to ID no. 76376, 1222.



interex

Shared Knowledge.

Shared Power.

As a not-for-profit association of HP computing professionals, Interex is dedicated to meeting the information, education, and advocacy needs of its members worldwide.

In today's world of rapidly changing technology, Interex puts hands-on solutions to hardware, software, and operating system problems at your fingertips. Because members actively contribute—exchanging ideas and sharing solutions—Interex is a vital link in the transfer of HP expertise.

Operating independently from Hewlett-Packard, Interex has more than 20 years of serving HP computing professionals. Through its publications, conferences, and volunteer committee structures, Interex has the qualifications to represent you, a valuable member of the HP user community.

Interex® is a trademark registered in the U.S. Patent and trademark office.

your most valuable asset

and how to cover it.

Corporations protect their people with comprehensive health insurance.

They protect their buildings and capital equipment with catastrophic insurance. But what about the asset that many professionals feel is their most valuable? Data.

We are Concorde Technologies. And our network backup solutions are really data insurance. Case in point? Our 2/10 DLT® Tape Systems. Perfect for midrange applications, these high performance, 200-gigabyte libraries house two DLT 4000 tape drives and up to 10 DLT tapes. Their dual access technology optimizes system speed and delivers "lights out" backup, even if a drive goes down. An intuitive front panel with a diagnostic display gives you the status of all important library operations, instantly. And because you can connect up to nine 2/10's together where all drives can access all tapes, Concorde Technologies gives you a solution that meets your needs today — and tomorrow.

Bottom line, we give you an edge that goes beyond any specific boxes we sell. That edge is knowledge.



We know heterogeneous network backup as no other HP Channel Partner does. We know the right process, products and service. We know how to analyze your requirements and design the optimal solutions. And we know how to install, train and support the system.



In short, when it comes to your company's data, your asset is covered with Concorde Technologies.

Solutions beyond the box.

800-359-0282 • (619) 536-5500 • Fax (619) 566-4396

email: sales@concordetech.com

URL: <http://www.concordetech.com>





Question & Answer

Q: I have a new optical autochanger and with so many surfaces, it looks like a sysadmin nightmare. Do you have any tips on how to manage these devices, especially creating device files and formatting?

A: A typical 40-GB changer will have 32 discs or 64 separate surfaces, which means 64 separate file systems. For operating systems 9.xx and earlier (that support the autochanger), the virtual mount capability allows many or all of the platters to be mounted at the same time. Access to a given platter is simply by reference to the mount point or files/directories within the mount point.

At 10.xx, this capability has been removed and only platters that have an available drive can be mounted. An alternative to this change in support is to use HP OmniStorage, which will manage disc surfaces in a more transparent manner.

When first installing the autochanger, the two drivers *autoch* and *autox0* must be gen'ed into the kernel. Once done, attaching the autochanger to the system and rebooting will cause *insf* (800 Series) to install the device files for the first 16 platters. To install the rest of the device files for all platforms, use SAM: select the Disks, File Systems, then Magneto-Optical and choose the top menu Actions selection to add the additional device files.

When initializing a large number of platters at the same time, you can speed up the process by taking advantage of autochangers that have more than one drive. The two scripts (odd and even platters, *Listing 1*) can be run at the same time on a changer with two drives. Both drives will have a disc installed automatically and the formatting will proceed simultaneously. If you have an autochanger with four drives, you can replicate the scripts and adjust each script to run a different platter at the same time.

Be careful when you do a lot of operations on several discs at the same time. Autochanger thrashing will occur if you try to format both sides of a disc at the same time. By doing all batch operations from a script, you can control which disc is being formatted or *newfs*'ed. The parameters in *Listing 1* for *newfs* are:

```
newfs -L -n -m 0 -i 8192 <device_file> <disktab_entry>
```

where:	-L	= use long filenames (255 chars vs. 14 chars)
	-n	= do not reserve a boot area at beginning
	-m 0	= minfree = 0
	-i 8192	= bytes per inode

Most HP-UX installations will need long file names—very little space is consumed with this option. *-n* will save some megs since the optical discs will most likely not be bootable discs.

For autochanger platters, using *minfree* = 0 is generally OK since the platters would never be used for critical file systems such as */usr*. Minfree reserves 10 percent of the disc (by default) so that after a disc mount point is full, 10 percent more will be available for root, ideally to log in and fix the full file system.

Continued on Page 10

**Quick: Name the RAID system picked #1
by PC Magazine, BYTE and LAN Magazine.**



StorageWorks™

When industry's toughest independent critics put RAID subsystems to the test, they came to a unanimous conclusion. StorageWorks™ RAID subsystems from Digital™ delivered outstanding availability, reliability, and economy—plus industry-leading performance. All designed as scalable solutions for today's multivendor PC and UNIX networks, to protect your investment from desktop to datacenter.

To find out more about Digital StorageWorks RAID solutions, call Workstation Technologies at (800) 663-1966.



Workstation
Technologies, Inc.

"Solutions for Your Open Systems Computing Environment"

1-800-663-1966

Outside the U.S.: 603-890-6700

Digital, the Digital logo, StorageWorks and the StorageWorks logo are trademarks of Digital Equipment Corporation. Other trademarks are property of their respective holders.

CIRCLE 20 ON READER SERVICE CARD

In the two scripts, you will see reference to *-i 8192*, which controls the amount of data associated with an inode. Decreasing this value (minimum is 1024) will increase the number of inodes, while increasing this number will decrease the inode count. In both cases, the amount of free space for data files will change with the inode size; larger *-i* parameters will result in more free space (see *Table 1*).

TABLE 1

The basic command is: `newfs -L -n -m 0`

Add the parameter *-i #####* based on the table below:

Disktab Model	Bytes per inode:	Bytes avail per surface	Inodes Avail per surface
hpS6300.650A	-i 1024	278.4 megs	266,108
hpS6300.650A	-i 2048	295.1 megs	142,780
hpS6300.650A	-i 4096	303.6 megs	74,492
hpS6300.650A	-i 8192	308.3 megs	37,244
hpS6300.650A	-i 16384	310.6 megs	18,620
hpS6300.650A	-i 32768	312.2 megs	6,204
C1716T_2X 0	-i 1024	564.3 megs	543,804
C1716T_2X 0	-i 2048	598.2 megs	291,644
C1716T_2X 0	-i 4096	616.8 megs	151,548
C1716T_2X 0	-i 8192	626.3 megs	75,772
C1716T_2X 0	-i 16384	631.0 megs	37,884
C1716T_2X 0	-i 32768	633.5 megs	18,940

To get 308 megs per surface on a 1X (600-meg platter), use a *-i* value of 8192. The above table is for 1024-sectored discs. The available bytes will be slightly lower for 512 sectored discs. Note the tradeoff between the number of inodes (essentially the same as files) and capacity.

The autochanger is an ideal CD-ROM server, especially for something like LaserROM or Application or Core CDs. By copying the entire CD-ROM to a surface (always use 1.2-GB media with 600 MB per side), you can share the discs on a network. For install/update, the usage is quite low and the chance that two or more users will require both sides of the same platter will be relatively low.

LISTING 1 *Platter Initialization***First script (odd platters)**

```
#!/bin/ksh
# Format odd numbered platters (slots), side A then side B
#
for PLATTER in 1 3 5 7 9 11 13 15 17 19
do
    for SURFACE in a b
    do
        echo Initializing Platter $PLATTER, Surface $SURFACE
        newfs -L -n -m 0 -i 8192 \
            /dev/rac/c0d${PLATTER}${SURFACE}s2 C1716T_2X
    done
done
```

Second script (even platters)

```
#!/bin/ksh
# Format even numbered platters (slots), side A then side B
#
for PLATTER in 2 4 6 8 10 12 14 16 18 20
do
    for SURFACE in a b
    do
        echo Initializing Platter $PLATTER, Surface $SURFACE
        newfs -L -n -m 0 -i 8192 \
            /dev/rac/c0d${PLATTER}${SURFACE}s2 C1716T_2X
    done
done
```

Q: How can I automate ftp to push/pull files from remote systems?

A: Here's an example of a batch job set up as a script. The use of the construct `<<EOD` is known as a 'here' document, which acts as if ftp were receiving the commands from stdin. The nifty part is that items within the `<<EOD` until the line `EOD` will be expanded by the shell so that parameters and *ENV* variables can be used. The special word *EOD* can be anything (such as `<<EnDoFmYdOc` or `<<DONE`) as long as the string does not appear in the enclosed text.


```
#!/bin/sh
#
# ftp to the remote system...
# $1 is the remote machine
# $2 is the logon
# $3 is the password
# $4 is the command (i.e., get or put)
# $5 is the source
# $6 is the destination
#
/usr/bin/ftp -v -n >> /tmp/batchftp.log 2>>
/tmp/batchftp.err <<EOD
open $1
user $2 $3
$4 $5 $6
quit
EOD
exit
```

This is a template, since you may wish to perform additional steps such as *chmod* or *chown*, or even create additional directories. You may wish to add *-i* (to turn off the prompt for multi-file tasks like *mput*). The results of the ftp task are sent to two files for logging: stdout will have the commands and responses while stderr will have error messages.

Q: I've messed up my I/O configuration on a Series 800 computer. How can I reset the configuration file?

A: The file */etc/ioconfig* contains a description of the I/O devices and associated drivers. When it becomes corrupted, *rmsf* and *insf* may not operate correctly. If *ioinit* cannot rebuild */etc/ioconfig*, it may be necessary to replace */etc/ioconfig* with the 'empty' version. There are two methods:

1. An empty copy is stored in */etc/newconfig*. Simply copy this file to */etc* as in:

```
cp /etc/newconfig/ioconfig /etc
```

2. Create an empty file using the vi editor:
(be sure that vi works correctly for the terminal being used)

```
mv /etc/ioconfig /etc/ioconfig.old
vi /etc/ioconfig
```

Type the following chars: *!"IO^V^Desc*

where: *i* means insert new characters

!"IO is the empty ioconfig requirement

^V means press CTRL and V at the same time

^D means press CTRL and D at the same time

esc means press the ESCAPE key

then save the file with: *wq!*

Now the computer can be rebooted with *shutdown -r 0* or *reboot -q* and the *ioconfig* file will be recreated. Note that device files will not be affected by this task so it is probably a good idea to do this task in single-user mode.

After the reboot starts, interrupt the boot process and go into single-user mode to make sure the device files still match the hardware addresses. If not, you can delete the device files using *rmsf -H* (don't worry about active disks—the device files can be removed and added on a running system). Then you can add the device files back again using *insf -H*, and to force a specific naming convention (i.e., Logical Unit or LU number), use the *-l* option. Note that existing LVM disks *must* keep the same device file names. These names are hardcoded into several config files and cannot be changed easily.

Q: I have a 735 running HP-UX 10.01. I've noticed lately that my application transactions have been taking longer. How do I start to analyze the problem?

A: Performance analysis is a very complex process. It is a service that HP offers on a consulting basis. The important thing to remember is that performance is very relative. There are very seldom definite thresholds to go by. Also, what one user finds acceptable performance another may not. Since the performance statistics are relative, the best thing for you to do is gather data when the system is performing well. As a result, you will have a "standard" for comparison.

There are three bottlenecks on a system: disk, CPU, and memory. Let's first discuss the tools that are available with the core operating system. *vmstat(1)* is a quick way of getting an overall picture of performance on the system. It shows such things as CPU activity, paging, and disk activity. It is executed on an interval and iteration basis. The following command will report statistics at 5-second intervals for a count of 10.

Continued on Page 12


```
# vmstat 5 10
```

sar(1) is new to the 700s at 10.x. It also is run with parameters representing interval in seconds and the number of iterations. *sar(1)* is a bit more robust than *vmstat(1)*. It has many different options that show such data as buffer activity and system call use. Another handy aspect of *sar(1)* is that it shows the size of the kernel process table and kernel file table. Please refer to the man page for *sar(1)* for more information.

Some other tools that are available are *top(1)*, *uptime(1)*, *iostat(1)*, and *ps(1)*.

There are also some tools that can be purchased as separate applications. The most popular is Glance. It is an extremely useful tool. Glance has the advantage of reporting a great deal of per process information. *top(1)* is the only tool that also reports per process, but it is much less verbose.

Now, let's discuss the bottlenecks.

CPU

The symptoms of this bottleneck are usually high CPU utilization and a high run queue. *vmstat(1)* and *sar(1)* are good tools to use to check those values. Utilization of 100 percent alone does not mean you have a bottleneck. That simply means you are fully utilizing the resource. However, the run queue shows runnable processes that are waiting. The acceptable number of waiting processes is one of those relative values discussed earlier.

What can you do besides purchase a more powerful CPU? If this problem has come about suddenly, you might try to look for a single process that is hogging the CPU. You can do this with *top(1)* or *ps(1)*. Find out from the appropriate support resource if this is "normal." If you have a process that you know will

be CPU intensive, you might try adjusting its system *nice* value. See the man page for *nice(1)* or *renice(1)*. This value comes into play when the process scheduler is recalculating process priority. It will make more of a difference for a short-running process. Another thing you can do is better organize the run schedule of CPU-intensive processes. If possible, delay some to less heavy times of the day with the use of *cron(1)*.

Disk

The symptom of this bottleneck is high disk activity. Often times you will also see low CPU utilization as processes are having to wait for I/O transactions to be completed. Again, you must have benchmark data for comparison. You can use *iostat(1)*, *sar(1)*, and *vmstat(1)* to check out your disk activity.

Your single most important goal is to spread the I/O contention evenly among your disk drives. Remember to include swap areas in the equation as well as your data areas. If some of your data areas are on file systems, you may want to talk to your vendor about file system parameters such as block size.

Memory

Most often a memory bottleneck will mask itself as a CPU bottleneck or a disk bottleneck. If processes don't have enough memory, the system will start to page. The paging daemon *vhand* will use CPU to do its work, and I/O to disk is going to increase as the pages are transferred to the swap areas. This is why it is very important to consider your swap configuration thoughtfully. *vmstat(1)* reports paging information.

The first thing to consider is the size of your buffer cache. At 10.x, the buffer cache can be a static size or vary between a configurable minimum and maximum. If you reduce the size of your buffer cache, you will have more memory for running processes, but your file system I/O might suffer. Another option is to find the memory hog process and "help" it finish faster. Again, you can use the system *nice* value, but this time you will make it less "nice" to the rest of the system. Another idea is to call your application vendor and inquire about application parameters that would affect memory utilization.

Q: We are using a new application on our 712 workstation running HP-UX 9.05. We are receiving the following error from time to time:

```
PID XXX was killed due to stack growth failure.
```

```
Possible causes: insufficient memory or swap space, or stack size
exceeded maxssiz.
```

A: The process either is running out of total swap or is exceeding the size of the process stack. The size of this stack is determined by the configurable kernel parameter *maxssiz*. This is a per process value. Its default value is 8 MB.

First, let's talk about the output from *swapinfo*.

Continued on Page 14

Why not the BEST?

The BEST Products
The BEST Performance
The BEST Warranty
The BEST Value

Unattended Backup Solutions



Tape Library Autochangers

- HP 4mm DAT, Exabyte 8mm, 3480, QIC, DLT, S-VHS, and AMPEX DST 19mm drives
- Up to 600 cartridges per changer

Autoreply

- Allows any backup program to control Autochangers (TAR, CPIO, DUMP, FBackup, etc.)

QuadraJet

- **Tape Array** controller
- 400% increase in speed and capacity
- Automatic tape mirroring


UNISTORE®

- Unattended backup and Hierarchical Storage Management (HSM) for CA-UNICENTER

Backup Programs for UNIX and Novell

- Omniback and Omniback-Turbo, Hi-Back, Networker, ARCserve, Budtool, DATtool, OpenVbackup, and others

Storage Solutions



Disk Drives and Towers

- Fixed or removable drives
- Fast and Fast-Wide options
- HP disk mechanisms with 5-year warranties

Memory Expansion

- For HP 1000, 3000 and 9000

RAMdisk

- For HP 1000 and 9000

Optical Disk Drives

- Rewritable, Write Once, and CD-Recordable



Optical Disk Autochangers

- 3-1/2", 5-1/4" and CD (Up to 144 cartridges per changer)

Custom Solutions

Custom Device Drivers

- For HP 1000 and 9000

Custom Programming

- Consulting and Programming Services

Robotic Control Software

- With API, Interactive, and GUI interfaces

Call with your requirements!

CIRCLE 8 ON READER SERVICE CARD

"I promise you—no service hassles, no matter who is servicing your HP system."

Rick Walsh
CEO



HERSTAL
AUTOMATION

7414 Manatee Street
Sarasota, FL 34243-1824 U.S.A.

Phone: 1-941-358-2001 • Fax: 1-941-358-2010


```
# swapinfo -tm
```

	Mb	Mb	Mb	PCT	START/	Mb		
TYPE	AVAIL	USED	FREE	USED	LIMIT	RESERVE	PRI	NAME
dev	48	4	44	8%	0	-	1	/dev/vg00/lvol2
dev	36	4	32	11%	0	-	1	/dev/vg00/lvol3
reserve	-	25	-25					
memory	17	11	6	65%				
total	101	44	57	44%	-	0	-	

The *-t* option is important as it produces a total line. This totals the amount of swap actually in use and the amount in reserve. (Processes will reserve swap when they start and if necessary as they continue to execute. The swap is not in use, but it is held for that process.)

As the problem process runs, watch the total PCT USED value. If it nears 100 percent, you might have a system swap problem. In that case, you will need to add more swap or reduce the amount being used on the system. That would entail removing processes or reconfiguring them.

If the total swap usage stays low, then you are likely running out of space for the process stack. This parameter can be increased via SAM or manually. Either way, it will necessitate a reboot of the machine. This will not cause processes to use more space for the stack; it will only *allow* them to do so. However, the default size is usually adequate. You might want to run this problem by your application support to ensure that improper execution flow is not the fault.

Q: I'm having a problem sending network mail from windows within VUE. Every time I type the "@" symbol, it deletes my address line.

A: The problem is that the kill or delete-line character is set to "@". Add the following line to *\$HOME/.vue/sessions/current/vue.resources*:

```
*ttyModes: erase ^H intr ^C kill ^U start ^Q stop ^S swtch ^@ susp ^Z
```

Or, from the Toolbox on the front panel: Toolbox->General->System_Admin->EditResources. That will put you in an editor. Add the previously mentioned line to the file, save it, and exit the editor. You must log out and back into VUE for it to take effect.

Q: I have a user bob who is a member of the group prod2. I have the following executable on the system.

```
-rwxr-x-- 1 jan prod2 16384 Aug 9 17:16 trans
```

bob cannot run this executable. Why not?

A: By default, a user only has the group privileges associated with his/her primary group. The primary group of a user is named in the user's entry in */etc/passwd*. You can check out the current group privileges with the *id(1)* command.

```
# id
uid=270(bob) gid=120(prod1)
```

It is possible to change the primary group affiliation with the *newgrp(1)* command. However, an alternative is to create a file called *etc/loggingroup*. You can refer to the *group(4)* man page. Usually, this file is a link to */etc/group*, but it can be a separate file with the same data format as */etc/group*. This file is used to define a default group access list for each user. It essentially gives a user privileges for his primary group and secondary groups at the same time. The list is built at login, so users will have to log back into the system to be affected by modifications.

```
# ln -s /etc/group /etc/loggingroup
```

(user bob logs back in)

```
# id
uid=270(bob) gid=120(prod1) groups=121(prod2)
```

Q: I have added a third-party networking application to my 712 running HP-UX 10.10. Now I am unable to generate a new kernel. The following message is displayed:

```
# mk_kernel -s system
Compiling conf.c...
Loading the kernel...
/usr/ccs/bin/ld: Unsatisfied symbols:
```

Why is this happening?

A: This can be a common occurrence when adding software that involves the kernel. The problem is that routines are being referenced

OUR 16th YEAR

QUALITY HP WORKSTATIONS - WITHOUT THE HP PRICE

RENT · LEASE · BUY

Immediate Delivery On Most Items

- 700 SERIES: Models 710, 715/50, 715/75, 715/100, 735, 735/125, 712/60, 712/80, 745i
- X-STATIONS: ENVIZEX, C270X
- All 300, 400T and 400E Series
- Memory, Features & Disc Upgrades for all Workstations
- CPU Upgrades: 715/50, 715/75, 735, 735/125, 425, 380, 360

*More Than 1500 Satisfied Customers
Throughout The USA and Worldwide*

For technical info, specs or pricing
call Mordy or Carol



C.S.U. Industries, Inc.

207 Rockaway Turnpike, Lawrence, NY 11559
(516) 239-4310 FAX (516) 239-8374

from object modules that are not being satisfied by the libraries in the link statement. One of two things has likely happened. Either the application libraries were not loaded correctly on the system or the libraries were not added to the make file properly due to problems with the tables in `/usr/conf/master.d`. This file, among other things, associates drivers with the appropriate libraries containing the driver object code.

First, make sure that you don't have a problem with the operating system itself. Edit `/stand/build/system` and temporarily remove any lines that refer to the application. Make sure you can build a kernel from this system file.

```
# mk_kernel -s system
```

If the kernel is generated successfully, you have a problem that is specific to the addition of the application software. Contact your application support representative. ■

General HP-UX questions are answered by Bill Hassell, a support engineer at the HP Atlanta Response Center. He can be contacted via e-mail at blh@hpuerca.atl.hp.com. Workstation questions are answered by Susan Potter, an HP-UX system support engineer in the Atlanta Response Center. Her e-mail address is sup@atl.hp.com.

Aliases, Baud Rates, and More

THIS MONTH I AM GOING to cover several small topics, all of which are important, but not sufficiently in-depth detailed enough to warrant a whole column.

Sendmail Aliases

In the September 1996 column I reviewed the sendmail 'Bat' book and gave an introduction to sendmail and its configuration. One area I did not touch on was aliasing. In sendmail, and most e-mail systems, an alias is used to send a message to one or more people, or to send one user's mail to another, without the sender knowing it.

There are several uses for sendmail aliases:

1. Forwarding root, uucp, adm, etc. e-mail to the "real" system administrator.
2. Creating a single alias for the entire company, department, or group for internal use.
3. Creating a single alias for the entire company, department, or group for external use.
4. Forwarding e-mail for an employee who is on vacation or who has left the company to another employee so it can be handled properly.
5. Defining multiple names for a user for external e-mail.

The first use is probably the most obvious. Since you should not be logging in as root as your primary user, it may be hours, days, or even weeks before someone checks root's mailbox. By aliasing root's mail to another user, that user will receive the mail and be able to act on it immediately.

The second use, setting up an alias for the company or a group, allows send-

mail to keep in one place a list of who belongs to a group, instead of each user having to keep his own alias databases. I have defined several aliases for my company (no, these are not the real ones, so don't try sending e-mail to them at *bwilab3.atl.ga.us!*): "all" for all employees, "sales_mkt" for the sales and marketing group, "techies" for the development staff, and "office" for the office staff. By using sendmail's alias feature, when an employee, like our latest co-op, begins work, I modify the "all" and "techies" aliases and the new employee automatically begins receiving e-mail sent to those groups.

The third use, an external alias, is useful for your customers to send e-mail to your Sales, Human Resources, or Technical Support employees without having to know a specific person's name.

The fourth use is important if you have certain users who receive e-mail that must be responded to or acted upon, because the sender cannot wait for the person to return from vacation. Also if any employee leaves the company, it may take some time for your customer base to know their new contact. You don't want to bounce their e-mail back.

The fifth use was something I did not have to do until recently. We hired an employee who has been in the industry for a long time. He had used the same user name for years, just changed his domain. This way, when he said "I now work for Bradley Ward Systems, Inc.," his contacts could send him e-mail at the same user name, but a different domain. Unfortunately, his long-time user name didn't match the corporate naming standard (which is a very, very complicated system: the user's first name!) (Yes we are that small, but not for much longer. ...) So I created an



LOOKING FOR A FASTER WAY TO LOAD YOUR DATABASE?

SyncSort UNIX can help you complete database loads, reorgs and reports in as little as *half the time*.

SyncSort combines high-speed sorting, versatile data manipulation features and the ability to handle a variety of data and file types. The result is a powerful, flexible tool for breaking database bottlenecks. To order SyncSort or for a free copy of our booklet, "Sorting and Relational Database Performance," please call or fax.

syncsort

Tel (201) 930-8200 dept. B6HUS
Fax (201) 930-8285 dept. B6HUS

alias on our mail server to accept his external name and map it to our internal name.

Setting Up an Alias

Setting up a sendmail alias is pretty easy: edit `/usr/lib/aliases` and add a line like the following:

```
alias: destination
```

where *alias* is the original address (such as root, or bob) and *destination* is the address the mail is to be forwarded to. The destination can be a comma-separated list of users. There can be only one alias in the file, but a user can appear in multiple destinations.

If you don't already have a `/usr/lib/aliases` file, copy the one in `/etc/newconfig/aliases` to `/usr/lib/aliases`.

Once the aliases have been entered, save the file and run `usr/lib/newaliases` to tell sendmail about the aliases.

The first thing about sendmail aliases you will note is that you must be root to edit the aliases file. So normal users can't change aliases or add their own. However, every e-mail client/viewer I have seen has the ability to add user aliases that behave the same way as sendmail aliases. The nice thing about using sendmail aliases is that they are global, so changes are made in one place, not in every user's personal aliases.

Modem Baud Rates

Two baud rates must be set when configuring a modem: the first is the rate between the modem and the calling/called modem. The second is the rate between the modem and the serial port and getty on the server. One of the most frustrating, and easy to fix, problems when setting up a modem is that the

getty on the server is looking for the modem at one baud rate while the modem is talking to the getty at a different rate.

The problem is caused by both sides "hunting" for the correct baud rate. Fixing the problem is easy if the modem is "intelligent" and almost impossible if not. By *intelligent* I mean that the baud rate between the modem and the serial port can be locked at a specific value. Usually this is referred to as "Baud Adjust." When Baud Adjust is enabled, the modem will try to detect the speed at which the serial port is communicating.

Unfortunately in most UNIX systems, the getty is trying to do the same thing. Look at `/etc/gettydefs`. Without getting into the gory details of what the fields mean, look at the first item in a "paragraph." It should be a number such as 1200, 2400, or 9600. Then look at the last line in the paragraph. It should look like:

```
#login: #2400
```

The number might be different. What this line does is tell the getty the next baud rate to try if the current entry cannot successfully connect to the serial port. This is fine for ASCII terminals that are hardwired to the server and always communicate at the same rate, but a modem connected to the serial port cannot guarantee a valid connection (or a specific speed each time). Or can it?

As I said earlier, most intelligent modems can perform a function known as Baud Adjust. This feature is nice for PCs running non-UNIX OSs but can cause havoc on UNIX systems. Why? By default the getty wants to hunt for the correct baud rate. If the modem is doing the same thing, the two may never sync

up. Think about it: The modem connects at 9600 and tries to connect to the getty. The last connection was at 2400 so the getty is still set at 2400. Since the rates are different, the user sends a break to wake up the server. The getty now goes to 4800, but the break sent the modem to 1200. The cycle will continue until the modem or the getty is reset.

So what to do? Lock down the baud rate of the connection between the modem and the getty. In `/etc/gettydefs` there is an entry for 19200 which automatically tries 19200 again when the getty receives a break. This way the getty is always at 19200. Edit `/etc/inittab` and change the baud rate at the end of your getty command to 19200. Execute a `telinit`, then kill the getty.

Next change your modem's configuration. Turn off Baud Adjust (usually \$BA) and force the Serial Port Baud Rate to 19200 (typically \$SB19200). Note, your modem may use different commands, or have a different maximum baud rate (older 2400's did). If your modem's maximum serial port baud rate is less than 19200, edit `/etc/gettydefs`, copy the 19200 entry (all four lines PLUS the blank line) and change the 19200 in all places to your maximum speed. Save the file and change your `inittab` file as described earlier.

Now what you have is a connection from your server to your modem, which is always a fixed rate regardless of how fast or slow the modem calling your modem can communicate.

Quickies

Want to change the first thing users telnetting to your system see? Edit `/etc/inetd.conf` and modify the telnet line as follows: instead of just `/etc/telnetd telnetd`, change it to `/etc/telnetd telnetd -b`

FILE_NAME where *FILE_NAME* is the name of the message you want them to see. Reload *inetd* with the following command: */etc/inetd -c*. Then telnet in to see how it looks!

Interested in public domain software for HP-UX? The Liverpool archives have a daily e-mail mailing list of the new or updated packages. Send e-mail to *Majordomo@csc.liv.ac.uk* and include "subscribe hpux-daily YOUR_E-MAIL_ADDRESS" in the body of the message.

YOUR_E-MAIL_ADDRESS is the address to which the daily update is to be sent.

In the May 1996 column I mentioned a site that contained some HP-UX 10.X release notes and manuals online. I received several messages telling me that the URL was not valid. In my last column I asked if anyone knew where to find this information. After sending the column off to my editor, I received some e-mail with a new URL that does work (or did in August): http://www.uiuc.edu/ccso/workstation_grp/hp_project.

That's it for this time. Keep the comments and ideas coming. ■

Chris Curtin, a software developer for Bradley Ward Systems, Inc. in Atlanta, Georgia, specializes in device driver development for factory automation on the HP 9000. He can be reached via e-mail at: chris@bwilab3.atl.ga.us.


Large Databases Require New Solutions for Backup !!!!!

FINALLY!

Backup 4 to 8 Gigabytes/Hour/Drive
Store 20 to 40 Gigabytes/Tape
Random Libraries from 5 to 264 Cartridges

DallasTools

Back-up, Restore and Archiving Software for Heterogeneous Networks



Dallastone also sells and supports all major brands –
Quantum, Breece Hill, ADIC and Odetics



DALLASTONE
2 Cote Lane
Bedford, NH 03110

Phone 603-647-8168
Fax 603-624-2466
Email dttool@delphi.com

CIRCLE 76 ON READER SERVICE CARD

Make the UNIX to MPE Connection

IX/92™

Full featured HP terminal emulation

New! Version 6

Faster File Transfer NS/VT Network Option Enhanced Script Language

Available for:


HP-UX Interactive UNIX SCO UNIX SunOS/Solaris

Software Licensing Corp., Suite 280, 930 Tahoe Blvd. Unit #802
Incline Village, NV 89451-9436

Phone: (800) 831-0882 or (702) 832-0881 Fax: (702) 832-0883

All trademarks are the property of their respective holders.

CIRCLE 180 ON READER SERVICE CARD



by David L. Totsch

Capturing Console Text to a File

FREQUENTLY, I FIND MYSELF wanting to capture text that is on the console into my log book or into a file on the system. It would be nice to have an inexpensive printer lying around to attach to the printer port on the console when needed, but the adage “what you don’t use you lose” applies to any computing equipment that just sits, including cheap little printers with desiccated ribbons and yellowed paper. Usually, these messages are system messages or Tomb Stones recorded by the Service Processor. Always, it is text I want to capture for posterity or further investigation (yes, laziness is the mother of invention—and computers are the culmination of a lazy person’s dreams). These nifty little consoles, the “smart” terminals that they are, have the information recorded in their memory. The difficulty is getting the information from the terminal’s memory into an ASCII file (without typing more than it would take simply to key in the message(s)).

Actually, it really is not much trouble at all to get that information out of the terminal’s memory and into a file. Merely get the information you want captured displayed. Then, from a login, run the *cat* command to capture standard input to a file. Use the Scroll Down key or shifted arrow keys to move back in memory until the text you want to capture is displayed. Now press the enter key (not the Return key—the enter key is on the far right of the space-bar row). Every line that you display by depressing the enter key will be recorded in the file you have *cat* redirected into. When you have displayed all of the lines you want to capture, tell *cat* to stop by entering ctrl-d. The file you were running *cat* to now contains the lines you displayed by pressing the enter key.

Here is a description of capturing the help information from the Access Port on a T500:

ctrl-b	get the system’s attention and enter the Access Port (you should see the “CM>” prompt)
he	display the AP help information
co	return to console mode
cat > myfile	capture to ‘myfile’
	press Scroll Down until the top line displays “CM> he”
	press Enter until you have captured all of the lines you want
ctrl-d	tell cat to stop capturing
clear	clear the screen (if you wish)
cat myfile	you should see the following:

```
CM> he
```

```
HP 28639A Access Port Revision 1.1 - 3245
```

```
CA - Configure system remote support modem port.
```

```
CO - Enter console mode.
```




New and Refurbished
1000/3000/9000 HP Equipment



NORCO
COMPUTER SYSTEMS, INC.

21337 Drake Road, Cleveland, OH 44136-6620
(216)-572-4040 (800)-892-1920 Fax (216)-572-0636
eMail: norcocomp@aol.com web: www.norcocomp.com

is your system trying to tell you something?

The total product & service package can only be found at NorCo.
Call today and see what we mean when we say,
"New or used, NorCo means a good deal and a good deal more!"



ISN'T IT TIME FOR NORCO?

CIRCLE 71 ON READER SERVICE CARD

- CS - Copy screen from local console to remote console.
- DI - Disconnect line to remote console terminal.
- DR - Disable access by a remote console terminal (enable session
access).
- DS - Disable display of system status line during console mode.
- ER - Enable access by a remote console terminal (disable session
access).
- ES - Enable display of system status line during console mode.
- HE - Display this screen.
- LR - Lock remote (disable modem access).
- RS - Use RS command in SP mode.
- SE - Transfer remote terminal from console/control to session mode.
- SP - Enter Service Processor mode.
- TA - Initiate Access Port selftest.
- TC - Use TC command in SP mode.
- TE - Send message between the local and remote console terminals.
- UR - Unlock remote (enable modem access).


CM> co

Now, the next time the console coughs up some of those cryptic little foreign messages and you find yourself meticulously recording them in your log book, make the system do the work for you. Of course, the underlying requirement is that the

system be operational. Sigh. You can always beg for an inexpensive printer to attach to the console...but that is another story. ■

After serving several different organizations over the past eight years as a system administrator with various flavors of UNIX, David L. Totsch still enjoys the profession. He also enjoys discussing UNIX with just about anyone. At present, he is working with HP-UX systems and wide-area networks for a Fortune 100 company in the Piedmont area of North Carolina. He can be reached via Internet: dtotsch@wfu.edu.

Protect Your Computer



As You Protect Your Business

Illustration by Michael Ehrhardt

by John A. Pezzano

Basic computer security is like basic security in any other area. The concepts are the same; only the specific implementation changes. Therefore, when researching information on how to explain to managers and administrators the basic steps to take to protect their computer systems, I didn't have to go to a technical library to define the concepts. I simply called my local police department's community relations unit and asked them for some brochures on protecting a small business from crime. The information they provided had all the steps needed. I just had to interpret them for the computer.

Physical Security Procedures

It is important that the perimeter of a computer system be protected from those who are trying to break in to it. The first line of defense against an outside attacker is making sure the basic entryways to a system are secure.

Make sure that all exterior doors are constructed of solid wood or metal for maximum security.¹

The doors of a computer system are its terminal connections. Solid doors require good passwords to keep a hacker from repeatedly beating against the system and breaking in through a weak password. Brute force methods, such as trying all sorts of common password combinations, are often used by the patient hacker. Good passwords, like solid doors, are hard to break.

Good passwords are at least six characters long and have at least one

non-alphabetic character. Common dictionary words, login names, or family names make for weak doors. Misspelled words (*mispeled*), words with embedded numeric characters (such as *h0uses* instead of *houses*), words containing non-alphanumeric characters (*mom&pop*—or better, *m0m&p0p*), first letters of phrases (*apsiape* for *A penny saved is a penny earned*) are all much more solid.

For additional protection, a bar across the exterior of the door, held in place by a padlock, is good security.²

Use the extra level of a dialin password for additional protection for those exposed modem doors. It constitutes one more layer of protection and slows down the hacker attempting to break in. Like the construction of the door itself, the dialin password bar should have a password that is difficult to break. The */etc/dialups* and the */etc/d/passwd* files, if they exist, require that an additional password be provided. This password is not user-dependent, but instead is based on the login shell program of the user. Since any port names specified in the *dialups* file require the use of the "dialin" password, this added protection can be used not only for modems but for terminal ports out on the factory floor, for example.

For maximum security, a dead bolt or jimmyproof rim lock with a one-inch bolt should always be used.³

For maximum protection, don't let the hacker even get past the modem. Dialback modems, which require addi-

tional passwords or dial back the caller at a known location, provide maximum protection. As an adjunct or alternative, there are security systems that require a special box that transmits a security key. Others have credit card sized displays that provide a key that changes every minute.

When the doors are not in use, disable them. Modems and/or local ports can be disabled or enabled by selectively using *init* states. For example, *init* state 2 might enable all ports. At night when dialin terminals might be disallowed, change to *init* state 3, which permits local ports but not dialin ones.

If you have a bank of modems, another possibility during off hours is to disable all but one high-security modem that has a separate number for administrator use only.

Too often administrators fail to consider this most basic step: If you turn off the modems when they are not in use, nobody can dial in to them no matter how sophisticated they are.

Lock unused doors.⁴

Doors that are not going to be used for a while should be securely locked so they cannot be opened except from inside. Disable passwords for users on extended absence (long vacations, out of country, working on special projects where computer access is not needed, etc.). Simply replacing the password in the */etc/passwd* file with an asterisk (*) results in that account being disabled as no password encrypts to this character.

Continued on Page 24

¹Commercial Security—a guide to Crime Prevention, Dekalb County, Georgia Department of Public Safety brochure

²ibid.

³ibid.

⁴How to Crimeproof Your Business, U.S. Department of Justice publication

*Skylights, ventilation ducts, and fire escapes may tempt a burglar because they're usually not visible from the street.*⁵

Alternate entrances to a system must be protected. Burglars don't always come through a door or window. They find any weak entrance, so look for them.

Alternate root user names should be avoided if at all possible. Although sometimes necessary, alternate root user names are a weak point. Often these IDs are buried in the password file and so are forgotten. When an administrator leaves, these alternate IDs may not be removed or have their passwords changed. Since an alternate root user is often set up for convenience, many times the password is also chosen for convenience and fails to meet the stricter guidelines set for other users.

Network capability that is not required should be disabled. If your system is, as most are, on a network, disable those features that are not required. If you use only telnet but not *ftp* or *rcp*, disable those capabilities by deleting the services in the */etc/services* file or, for HP-UX, limiting access via the *inetd.conf* file.

A common method of becoming the root user is to find suid programs or scripts. Suid programs give the capability of the program owner to the program user. So, for example, while a user cannot edit the *passwd* file, he can change his own password by using the *passwd* program, which gives him root capability only for doing this one job.

Suid (set user id) programs run using the capabilities of the program owner rather than those of the user. Programs that must run suid should be known and checked. Suid programs must not allow ordinary users to write into them as this would allow a user to change their operation.

Suid scripts are even more vulnerable and should be avoided if at all possible.

To protect against forgotten or unknown suid programs and scripts, disks can be mounted so that the suid bit is ignored even if present. By limiting suid programs to one or two disk volumes and mounting the others without suid capability, even a forgotten potential entryway can be blocked.

*Some businesses may need to install an alarm.*⁶

When someone breaks into a system, an "alarm" should go off. Such alarms can be triggered by commercial security programs or public domain ones such as *Tripwire*, which checks for programs that have changed. In addition, the administrator can write simple scripts that look for unusual activity and report on it. For example, a script might check for too many failed login attempts from a port or Internet address or too many different user names trying to log in to a single serial port or modem. It is important that any security program not provide so much information that essential data gets lost and missed. For example, if there are 100 root *sus* a day, a summary report should filter out the expected ones such as *root->root*, *cron*

jobs, standard administrator *user->root*, etc., and display only those that are unexpected or that occur during odd hours. Otherwise, finding the unusual entries can be impossible.

*For an expert appraisal of your security needs, ask for a premises security survey by your local law enforcement agency, or check with a reputable professional consultant.*⁷

There are numerous security programs that you can run to do a security survey. *C.O.P.S.* (Computer Oracle and Password System) is a good general-purpose public domain security checker. It will test for weak passwords, suid programs, the ability to break through known weaknesses, and more.

SATAN is a public domain system administrator's tool for checking network security.

Crack is a public domain tool to do extensive testing of passwords.

Third-party software packages are also available. The main advantage is that their reporting is more readable and understandable and support is available from the vendor.

An outside security consultant can give you an independent look at your security weaknesses and recommend improvements. Hiring such a consultant *before* your system is successfully attacked can save a lot of time, money, and heartache. Always verify the credentials and integrity of any outside security consultant as you are exposing all your weaknesses to this person.


Continued on Page 26

⁵ibid.

⁶ibid.

⁷ibid.

When you run your Data Center right, you are a winner!



DESIGN
Plus
3000
INCORPORATED

When your Data Center has problems it's a race against time until they are corrected. We have JMS to schedule and control your batch jobs, and CALLBACK to monitor and notify you of critical conditions. You will be notified by phone, pager or alpha pager of any user defined conditions or other problems that occur within your Data Center.

JMS JOB MANAGEMENT SYSTEM

JMS puts you in control of your Data Center with scheduling and execution of batch jobs. JMS has no command language to learn so it's easy to use and yet has sophisticated scheduling capabilities that are second to none. JMS network capabilities allow scheduling and job dependencies across the network. Restriction features allow you to limit those pesky system hogs. Ad hoc streamed jobs may be incorporated to allow complete batch job control.

CALLBACK SYSTEM DOWN DETECTOR AND JOB NOTIFICATION SYSTEM

Voice/Beeper/Digital or our new Alpha Pager notification for:

- Adverse Temperature • Power Loss • System Hangs • Job Aborts • Printer Status • Physical Conditions • User Requests • Logon Security • Reply Pending • Console Messages • Special Job Events • Spoolfile Scanning • And much more...

We are the leaders in tools for the "Lights Out Environment." Call us today for a free 30 day demo. MPE/iX or HP-UX.

Corporate Headquarters

3470 Pipebend Pl. NE • Suite 120 • P.O. Box 13086 • Salem, OR 97309-1086
Phone: (503) 585-0512 • FAX: (503) 585-1706 • E mail: Design3000@aol.com

International Sales Office

System Software Intl. • Oakmoore Court • Kingswood Road • Hampton Lovett
Droitwich, WR9 0QH, UK • Phone: +44 (0) 1905 794646 • FAX: +44 (0) 1905 794505



Internal Security Procedures

If someone penetrates your outer perimeter, you want to make it as difficult as possible to do anything destructive or to steal your data.

*Use an Underwriters Laboratory listed money safe, bolted to the floor...*⁸

When you need higher protection, you should turn on "Trusted Systems." This provides U.S. Government compliant C2-level security. Trusted Systems optionally allows auditing to a high level of detail, although at the cost of system performance. At its most basic level, Trusted Systems provides for a secured password file. In ordinary UNIX, the encrypted passwords are visible to everyone. A prime method of attack is to gain access to a system as an ordinary user, then copy the password file and, using tools such as *Crack*, attempt to break the passwords offline. By providing a publicly viewable password file that shows no passwords and a secure one that does, this method of attack can be thwarted.

For a much higher level of security, the compartmentalized B1-level UNIX provides for mandatory access control of files and separates the traditional root user capability into separate independent functional areas. Thus root might be given the capability to do everything except auditing and the audit user can check to see whether root is doing something inappropriate.

*Avoid keeping large amounts of cash in view.*⁹

You don't want to provide an entice-

ment to outsiders. The */etc/issue* file should not contain anything that welcomes users. It should also avoid providing information that will identify the type, manufacturer, or revision of the system and should not provide any information on the company. Only authorized users should be able to get to the *issue* file. A simple trick is to have the login prompt (defined in */etc/inittab*) changed for modem ports so that the system appears not to be a UNIX machine. Those familiar with UNIX will try known UNIX login names such as "root" and "quest."

Many businesses keep small amounts of ready marked cash in case of a robbery. In a similar vein, the "quest" account can be set up so that if anyone attempts to log in with that ID, an alarm is automatically triggered.

*Signs deter the shoplifter. They may state "Shoplifters will be Prosecuted."*¹⁰

The */etc/motd* (message of the day) or similar file should make it clear that the system is for authorized users and that violators will be prosecuted. It should also make the point that users understand that their actions on the system may be monitored. Businesses should consult with a knowledgeable attorney, if necessary, to help draft the correct wording.

A sample file with explanation is available from the Computer Emergency Response Team (CERT), operated by Carnegie Mellon University for the U.S. Government. CERT can be reached at cert@cert.org.

*Develop clear policies regarding security and theft.*¹¹

A written Security Policy is an absolute must for any organization. It should clearly state authorized and unauthorized activities and the penalties for violations of the policies. Issues such as what constitutes acceptable personal use, viewing of other users' files, downloading graphics files of questionable taste, hacking other systems, etc., must all be addressed. Help in writing security policies is available from reputable consultants. There are also commercial programs that will create customized policies.

The Security Policy should also include an *Incident Plan* that describes what action should be taken in case of a break-in. The appropriate sections of this plan must be readily available and understood by operators and others in case an incident occurs that requires attention before the administrator can be contacted. It should include the conditions and steps necessary to protect the system should a security incident occur.

*Include crime prevention information in your staff memos and newsletters.*¹²

Security briefings should be held on a regular basis. Company or organizational newsletters should include information on proper computer security. CERT advisories can be obtained from cert.org and such information should be distributed to all administrators. For Hewlett-Packard systems, HP provides the ability for customers to be on a mailing list for security bulletins and advisories.

⁸Protect Your Profits: Security Tips For Small Businesses, National Crime Prevention Council brochure

⁹Commercial Security—A Guide to Crime Prevention, Dekalb County, Department of Public Safety brochure

¹⁰ibid.

¹¹Protect Your Profits: Security Tips for Small Businesses, National Crime Prevention Council brochure

¹²ibid.

Train employees in proper cash handling and security measures.¹³

All users should be required to change passwords on a regular basis; the policy should be enforced by password aging. Users should be trained in the importance of security to the organization.

A default *umask* parameter for all users should be placed in */etc/profile* or */etc/csh.login* to provide default security that at the least denies write access to a particular file unless specifically overridden.

File security should be understood and used by all users.

Alert employees may be your best defense.¹⁴

Employees should understand what the likely signs of a break-in are and should learn to watch for bizarre occurrences, unusual activity, and strange changes on the system. Is someone logged in who is known to be on vacation? Is there activity on certain ports at a time of the day when there usually is none? Have some seemingly innocuous files been changed for no apparent reason? Do dialout telephone line charges seem unusually high or are they for destinations that don't make sense? Are there unknown files on the system? Is there something that just doesn't look right?

Protect Against Employee Problems

Your threats are not always from the outside. You have to protect against legitimate users committing malicious acts and administrators misusing the system.

Good key control is vital to any office manager or business owner. Make a "key list" of all individuals who have keys to the office or building.¹⁵

The number of users who know the root password should be limited. If certain users require root privileges for specific activities, then only those users should be given added capability and only for the specific programs. Both commercial products as well as public domain programs provide such capability. In addition, standard UNIX access control lists (ACLs) and more advanced user-specific ACLs such as those available on HP-UX should be utilized.

The root password, which is the master key, should be well protected and changed as often as necessary. Some organizations are afraid that the primary administrators may be unavailable when the root password is unexpectedly needed and so the password is given to extra people. An alternative is to keep the password locked in a high-level manager's safe and require contacting that manager or the manager's secretary to obtain it. Thus the password is available but someone outside the department now knows when it has been given out.

Revision control systems provide a method of checking that key application programs are not changed, or if they are, help determine what changes were made.

Require all terminated employees to hand in keys and have all locks rekeyed at least once each year.¹⁶

Passwords should be changed imme-

diately whenever an administrator leaves. Password Aging should be implemented to require that all users change their password on a regular basis.

Changing modem numbers can eliminate the knowledge that ex-employees might have about how to get into the system.

Maintain a log to record removal and return [of keys].¹⁷

Log system console output to a printer and avoid using the console for other purposes so that you have a log of key events.

Printing the output of the *su*log file provides a record of all attempts, successful or not, at becoming the root user.

Your best defense is frequent and thorough inventory control.¹⁸

Backing up frequently is the primary method of "inventorying" your files. System auditing adds another level of security as does turning on accounting. Auditing, which is used in Trusted Systems, can cause significant overhead, depending on the extent of auditing activated.

On HP-UX, you can audit your system files with *pdfchk(1M)* for 9.X systems and *suverify(1M)* for 10.X systems. Both of these provide a check of ownership, permission, size, checksum, links, and revision level.

It is vital that your system administrator do a regular audit of the *passwd* and *group* files to remove users who no longer use the system, and run COPS to do a regular checkup of the system.

Continued on Page 28

¹³ibid.

¹⁴*Commercial Security—A Guide to Crime Prevention*, Dekalb County, Department of Public Safety brochure

¹⁵ibid.

¹⁶ibid.

¹⁷*How to Crimeproof Your Business*, U.S. Department of Justice publication

¹⁸ibid.

Insurance and Enforcement

Insurance won't prevent crime, but it can help reduce your losses.¹⁹

There are two types of insurance to be concerned about. One involves system procedures to protect yourself, such as good off-site storage of backups, reliable and competent administrators, and loyal employees. In addition, the normal business insurance that protects against lawsuits and losses should be reviewed to see that it covers the following:

- computer break-ins
- employee destruction of system files
- loss of data and productivity because of malicious internal or external users
- liability for privacy violations or release of system information due to a break-in or user action
- liability for users (both authorized and unauthorized) hacking into outside systems from within your system

You need to know how to protect yourself financially and your legal obligations and responsibilities in case of a security incident. Senior management should confer with legal counsel for advice.

Burglary insurance requires certain security precautions.²⁰

You have a certain obligation to protect your business. If you do not, the

insurance company might believe that you failed to practice due care and might not be willing to insure a loss. Similarly, failing to take basic security precautions may be considered a lack of due care.

If you discover a burglary has occurred in your place of business, do not disturb anything.²¹

Should a break-in or loss occur, it is important that you disturb as little as possible in order for someone to determine what has happened. Date/time stamps on files will be changed by your recovering, viewing, or fixing files, and thus valuable information can be lost. You should take steps to protect the system and to maintain strict accountability for backups.

You should know what constitutes legal evidence in case of a break-in before you have one. Do you know when computer information is considered "business records"?

Cooperate fully with law enforcement and prosecutors.²²

Often companies fail to report or prosecute computer crimes, particularly when an inside user is involved. The result is that those criminals are free to prey on other companies. You should be prepared to provide the police with the information you have on what happened and when. You should learn how to gather such information or get a consultant who can provide it.

The boss who takes merchandise and office supplies without paying encourages employees to do the same.²³

You should practice good security yourself. This means that your passwords should meet strict criteria. You should not be using the computer for personal use not specifically authorized by your security policy for all users. If you use the computer to browse the Net and download files of questionable taste, so will your employees.

The basic goal that you should pursue is to educate yourself and your employees in the actions and behaviors of most shoplifters.²⁴

You should understand how hackers break into systems and how they compromise security so you can recognize the signs of a break-in and take steps to protect yourself.

Learn about crime in your neighborhood and what is being done about it.²⁵

Read books and articles on computer security and pass them on to your users so they know how to protect themselves and your system. There are numerous books available in computer stores and copies of the U.S. Government's Rainbow Series are available from the National Security Agency, Fort Meade, Maryland, (410) 766-8729. Even many local law enforcement agencies can provide information on protecting your computer.

¹⁹ibid.

²⁰ibid.

²¹Commercial Security—A Guide to Crime Prevention, Dekalb County, Department of Public Safety brochure

²²How to Crimeproof Your Business, U.S. Department of Justice publication

²³ibid.

²⁴Commercial Security—A Guide to Crime Prevention, Dekalb County, Department of Public Safety brochure

²⁵Protect Your Profits: Security Tips for Small Businesses, National Crime Prevention Council brochure

There are security sources on the World Wide Web including those from CERT and the U.S. National Institute of Standards and Technology at <http://csrc.ncsl.nist.gov>.

Summary

Security concepts are the same whether you are protecting a small retail store or a large corporate computer. Good security policy, a well-protected perimeter, good internal security practices, and knowledgeable action in case of a problem all contribute to a less vulnerable computing environment. ■

John Pezzano has been supporting the HP-UX operating systems for Hewlett-Packard for over ten years. He is the author of numerous articles on a wide variety of topics published in hp-ux/usr and Interact magazines from Interex. Pezzano currently works for HP in Atlanta where he is a member of the North American Escalation team supporting HP-UX. He specializes in serial data communications and system security.

Hewlett-Packard 9000

It's Our Specialty

200/300 Series	400 Series	700 Series	800 Series
216/236/217	425e	705/710	E/F/G
310/320/330	425t	715/720	H/I/K
350/360/370	425s	730/750	
318/319/340	433s	735/755	
345/375/380		C/J Series	

We also carry memory and interface for all of our workstations

Printers	Mass Storage	Plotters
2225A/B/C/D	9121/9122	DesignJets
3630A PaintJet	9153A/B/C	DraftPro
C1602A PaintJet XL	C2254HA	DraftMaster
LaserJet II/IIiD/IIIP	C2440HA/JA	Desktops
LaserJet III/IIID	C3232A	Electrostatic
LaserJet IIIP/IIIsi		
LaserJet 4L/4P/4,4+/4SI/4V/4MV	C2213A/D	
	DAT Drives	
	CD ROMs	
	Optical Drives	

Specials
DesignJet 650C Plotter



*We offer large discounts,
outstanding service
and immediate delivery.*

TED DASHER & ASSOCIATES
 PH: 800-638-4833
 FAX: 205-591-1108
 E-mail: sales@dasher.com

CIRCLE 49 ON READER SERVICE CARD

Do You Know Where Your Security Holes Are? Find Them with SecurityAudit/UX!

**Have You Heard that UNIX is Notorious for Its Lack of Security Features?
Do You Know Where to Check to See if Your HP-UX System is Secure?
Do You Have the Time to Do This Checking Regularly?
Use SecurityAudit/UX To Do It All!**

EVEN IF YOUR SYSTEM IS SET UP CORRECTLY (AND HOW WOULD YOU KNOW IF IT WAS?) AND HAS NO SECURITY LOOPHOLES, IT CAN BE VERY DIFFICULT TO MONITOR SYSTEM CHANGES, AND TO ENSURE THAT SECURITY ISN'T COMPROMISED. THE MAGNITUDE OF THE PROBLEM INCREASES AS THE TOTAL NUMBER OF USERS CONFIGURED AND THE TOTAL NUMBER OF FILES GROWS. IT'S EASY FOR ORDINARY USERS TO CHANGE THE SECURITY OF THEIR OWN FILES TO ALLOW OTHERS TO ACCESS THE CONTENTS. A LOOPHOLE LEFT BEHIND INADVERTENTLY OR ON PURPOSE MAY BE EXPLOITED BY A DISGRUNTLED EMPLOYEE OR A HACKER TO BREAK SYSTEM SECURITY, SOMETIMES MUCH LATER.

SecurityAudit/UX PRODUCES OVER 40 REPORTS, CONTAINING DETAILED INFORMATION ON THE FOLLOWING CLASSES OF PROBLEMS:

- **User and Group-related problems, including weak passwords and non-unique identification numbers.**
- **File-system related problems, including historical tracking of files and detection of potential Trojan horses.**
- **PDF-related security problems, extended to detect changes in ACL specifications.**
- **Logging subsystems status display, and logfile analysis.**
- **Network-related status display and configuration weaknesses.**

SecurityAudit/UX RUNS ON ALL HP-UX BASED 9000 SERIES 700 AND 800 SYSTEMS, AND HAS BEEN SPECIFICALLY TAILORED TO ADDRESS PECULIARITIES OF HP-UX, SUCH AS PDF, ACL AND HP'S SHADOW PASSWORDS.

Call EUGENE VOLOKH for more info!



1135 S. Beverly Drive
Los Angeles, CA 90035 U.S.A.
FAX (310) 785-9566

CIRCLE 40 ON READER SERVICE CARD

The SpoolKick Procedure

by Bill Hassell

Every once in a while, the lp spooler subsystem will appear to lock up or the spooler will consume large amounts of CPU time without printing anything. Most of the lockup problems were solved with the 9.0 release, but occasionally a problem may develop. This article presents the series of steps to take to get your spooler back on line. The first section covers the procedure for HP-UX 10.x. If you are on Version 9.x, skip to page 32 for the section on the procedure for HP-UX 9.x.

HP-UX 10.x

1. Do an *lpshut* command. This requires root and the command is located in */usr/sbin/lpshut*. Now check to see if all

lp processes have terminated with

```
ps -ef | grep lp
```

If you see processes such as *lpsched* still running, kill these processes with *kill -9*, using the process ID number. These *lpsched*'s are associated with printers (or scripts) that are having problems, so you may wish to note which printers are affected. The parameters passed to the model scripts will indicate the printer name.

Change directories to */var/spool/lp* and remove the two files at the beginning of the list: *SCHEDLOCK* and *FIFO*. If they exist, remove any files with the word "lock" in them, e.g., *lpd.lock*, *tmplock*, etc., and the file named *Toutputq*.

2. Once all lp processes are gone, see that all printers are online and ready to print. If a printer is temporarily off line due to a paper jam or out of paper, use the *disable* command to stop the spooler from printing to that device.

Now type *lpsched* to start the spooler again. If the *Toutputq* deleted in step 1 reappears in */var/spool/lp* when you restart the scheduler, remove it again and clear *outputq* (*> /var/spool/lp/outputq*). Type *lpsched* again to start the spooler once more. Check to see that printing is working and that nothing is hung. If all is well, the spooler is now fixed. Check on any printers that were associated with *lpsched*'s that had to be killed manually. You may wish to start lp spooler logging to monitor spooler activity with:

```
/usr/sbin/lpsched -v -a
```

This will log printer activities into */var/adm/lp/log* and */var/adm/lp/lpana.log*. */var/adm/lp/log* can be read with just a *cat* or *more* command, while reading */var/adm/lp/lpana.log* requires the */usr/sbin/lpana* command.

3. If the spooler still locks up, you can use SAM to help fix the problem. A copy of the spooler control files is made whenever changes are made to the spooler using SAM. This means that the spooler can be recovered to this state using SAM's Previously Saved State option.

Once this is performed, start the spooler and check to see that all is well. Again, if problems develop, check the names of the printers that seem to be hanging the spooler to see if the printers are online. If not, use the *disable* command to take them out of active spooler status.

Also check the log file in */var/adm/lp/log*. There may be some information about the spooler problems logged there. The JetDirect software always logs activities in this file and network activities are typically logged into the file */var/adm/syslog/syslog.log* (but check the printer/plotter selftest page to see if a *syslog* server has been established).

4. If the spooler locks up again, the data in the *pstatus* file is probably corrupt. Perform step 1 again. Now we need to make a list of all the printers currently in the system. Use *lpstat -v* to get this list, and make a permanent copy by redirecting it into a file as in

```
lpstat -v > /tmp/lp.list.
```

You'll also have a list of printers that are remote or located on other systems. The unprinted jobs in the spooler are located in the directory

```
/var/spool/lp/request
```

where each printer has a directory of print jobs that have not been completed. Each job creates a pair of files: a short control file and the actual data for printing. If none of the print jobs are important, these files can be removed (details in the next step). Otherwise, the individual data files can be moved to a temporary directory. Since many files may have been printed with special options, you should also move the control files to be used later in reprinting the jobs.

Also note: If you have made customized model scripts in the directory */etc/lp/interface*, make copies of these custom scripts in the directory */usr/lib/lp/model* (the template directory). You may wish to rename the scripts in the interface directory since they will have the name of a printer.

Special note for JetDirect-connected printers and plotters: The model scripts for these printers are in a directory one level below the interface directory called *model.orig*. Be sure to save any customized scripts found in that directory too.

Also, the spoolkick procedure does not affect the bootp services used by JetDirect cards in this system; bootp is independent of the lp spooler.

5. With the spooler stopped, use *lpadmin -x* to remove every printer. If you have more than a few, you may wish to create a script to read the names of all the printers and run the *lpadmin -x* command automatically.

Once every printer has been removed, we need to move to the directory */var/spool/lp* to check on the remaining files and directories. Here is a summary of action steps:

- Remove the files *SCHEDLOCK* and *FIFO*.
- Clean the requests with *rm -r /var/spool/lp/request/**. Be sure to specify the */** at the end so that *request* will not be deleted.

Continued on Page 32

- The following directories must exist but have **NO** files in them:

receive

request

- The file *seqfile* can be removed or you can edit it to contain an ASCII number as the starting point for the next print request ID. (As a note, *seqfile* can be changed at other times but only when the spooler is shut down and there are no pending jobs to print.)

- The file *default* should be zero-length.

- Zero out the status files with:

```
> pstatus (or you can use cat /dev/null > pstatus)
```

```
> qstatus (or cat /dev/null > qstatus)
```

```
> outputq (or cat /dev/null > outputq)
```

Do not remove these files. If they are removed, they must be recreated as zero-length files with the correct permissions and ownerships.

- Now change directory to */etc/lp*

- The following directories must exist but have **NO** files in them:

cinterface

class

info

interface

member

sinterface

- Finally, we must verify that the existence, links, ownerships, and permissions are correct. At HP-UX 10.0 and higher, there is a very useful command: *swverify*. With this command, you can verify the entire lp spooler file set. Just use:

```
swverify PrinterMgmt.LP-SPOOL
```

If you get the message “Analysis Phase had errors,” check the file

```
/var/adm/sw/swagent.log.
```

It will report any missing files or incorrect ownerships and permissions, which you will have to fix. Note: It will also run *swverify* on any dependent products and mark them with a “+” symbol.

Once *swverify* reports “Analysis Phase succeeded,” we have completed the cleanup of the spooler.

Note: If you are using *JetAdmin* software to print to network printers, you will also need to run:

```
swverify HPNP
```

Again, fix any incorrect permissions or missing file problems until *swverify* reports “Analysis Phase succeeded.”

To verify that all is well, use */usr/sbin/lpstat -t*

```
/usr/sbin/lpstat -t
```

```
scheduler is running
```

```
no system default destination
```

At this point, we can add the first printer, preferably one that is close by so we can verify operation immediately. Start by testing a directly connected printer (either serial or parallel). You can use SAM to add the printer or the commands:

```
/usr/sbin/lpshut
```

```
/usr/sbin/lpadmin -p<prn> -v<devfile> -m<model>
```

```
/usr/sbin/accept <prn>
```

```
/usr/bin/enable <prn>
```

where: *<prn>* is a name for the printer, *<devfile>* is the device file, and *<model>* is the model script to be used from the choices in */usr/lib/lp/model*. Then, test the printer with:

```
lp -d<printer-name> /etc/group
```

which will produce a one-page (or so) listing of the *group* file.

This should completely restore operation of the lp spooler. Be sure to check on the latest patches. As of this writing, the latest patches are:

PHCO_7102 10.00, 10.01 (cumulative lp spooler patch)

PHCO_6632 10.10 (cumulative lp spooler patch)

You can obtain these patches from the HP Response Center, or through electronic means from HP Support Line. Connect to the Web service through

<http://support.mayfield.hp.com>.

HP-UX 9.x

This section describes the SpoolKick Procedure for 9.xx and earlier.

Continued on Page 34

HP RENTALS • HP RENTALS • HP RENTALS • HP RENTALS • HP RENTALS

HP WORKSTATION SALES & RENTALS



WORKSTATIONS

807	E55	NEW!	755	750	CRX	433 S	CRX	382	360	332
827	F20	J200	735	730	CRX-24	425 T	VRX	362	350	330
847	G30	K200	725	720	CRX-24Z	425 S	GRX	380	345	320
E25	H40	K400	715	710	CRX-48Z	425 E	PVRX	375	340	310
E35	I70	T500	712	705	GRX	400 S/T	EVX	370		

AND MORE . . .

Printers
Plotters

X Terminals
APOLLO ON Series

PC's
Memory

Discs
Test Equipment

800 Series
Data Acquisition

TSA is THE place for daily, weekly, monthly or long term rentals of Hewlett-Packard equipment.
Ask about our 6-12 month purchase plans. Equipment available for same day shipment.

1-800-422-4872

THE RIGHT EQUIPMENT. RIGHT NOW.



713/935-1500 • Fax 713/935-1555
Email: Info@tsa.com



2040 West Sam Houston Parkway N. • Houston, Texas 77043

HP RENTALS • HP RENTALS • HP RENTALS • HP RENTALS • HP RENTALS

New kids on the block...

“I bet there’s a whole bunch of ads in this magazine for **tape libraries**. But you won’t want to miss what my dad has to say, ’cuz his SL-400’s really neat—it’s easy to use, it won’t break, and you can even buy it from him, too.”



CIRCLE 103 ON READER SERVICE CARD

the spoolkick procedure

1. Do an *lpshut* command. This requires root and the command is located in */usr/lib/lpshut*. Now check to see if all lp processes have terminated with

```
ps -ef | grep lp
```

If you see processes such as *lpshed* still running, kill these processes with *kill -9*, using the process ID number. These *lpshed*'s are associated with printers (or scripts) that are having problems, so you may wish to note which printers are affected. The parameters passed to the model scripts will indicate the printer name.

Change directories to */usr/spool/lp* and remove the two files at the beginning of the list: *SCHEDLOCK* and *FIFO*.

2. Once all lp processes are gone, see that all printers are online and ready to print. If a printer is temporarily off line, use the *disable* command to stop the spooler from printing to that device.

Now type *lpshed* to start the spooler again. Check to see that printing is working and that nothing is hung. If all is well, the spooler is now fixed. Check on any printers that were associated with *lpshed*'s that had to be killed man-

ually. You may wish to start lp spooler logging to monitor spooler activity with

```
/usr/lib/lpsched -v -a
```

This will log printer activities into */usr/spool/lp/log* and */usr/spool/lp/lpana.log*. */usr/spool/lp/log* can be read with just a *cat* or *more* command, while reading */usr/spool/lp/lpana.log* requires the */usr/lib/lpana* command.

3. If the spooler still locks up, and you are running HP-UX Version 9.0x, you can use SAM to help fix the problem. If you are running 8.0x or earlier, skip this step.

At 9.0x, a copy of the spooler control files is made whenever changes are made to the spooler using SAM. This means that the spooler can be recovered to this state using SAM's Previously Saved State option.

Once this is performed, start the spooler and check again. If problems develop, check the names of the printers that seem to be hanging the spooler to see if the printers are online. If not, use the *disable* command to take them out of active spooler status.

Also check the log file in */usr/spool/lp/log*. There may be

...with seven years of experience

and a customer list that includes IBM, Disney, Shell, Chrysler, Price Waterhouse, Sallie Mae, and others.

Here's our story. For two decades, **IGM Communications** made tape access devices for the broadcasting industry. With great design and leading-edge technology, IGM captured over 80% of its market.

Seven years ago, IGM transferred its expertise in tape handling to the computer data storage market. IGM's ATL-5000 was the industry's first 8mm tape library. It's still being used in a wide range of applications—from storing animated drawings to collecting seismic

data—winning customer loyalty for its exceptional reliability.

In 1994, I bought IGM, launched a new division, **Straightline**, and began an aggressive strategy to produce robotic tape libraries. Our unique solution to data backup combines **4mm DAT technology** with easy operation, proven dependability, and manufacturer-direct sales and support.

So you see, we might be the new kids on *this* block, but we've been living in your neighborhood for years.



Get to know us better
by calling
206.865.8314



Software Compatibility:

UNIX Environments: Legato Networker, IBM Adstar, Raxco, Spectralogic, OSM, Tivoli; Netware Environments: Cheyenne ARCserve, Palindrome; NT Environments: Arcado Backup Exec

©1996 Straightline. All rights reserved. Straightline is a division of IGM Communications. All brand and product names are property of their respective holders.

677 120th Ave. NE, Suite 146 • Bellevue, Washington 98005
www.igm.com • e-mail: igm@igm.com

some information about the spooler problems logged there. The JetDirect software always logs activities in this file and network activities are typically logged into the file `/usr/adm/syslog` (but check the printer/plotter selftest page to see if a `syslog` server has been established).

4. If the spooler locks up again, there may be corruption in the `outputq` file. Stop the spooler with `/usr/lib/lpshut`, check that all copies of `lpsched` are gone (`ps -ef | grep lpsched`), and then zero the `outputq` file with

```
cat /dev/null > /usr/spool/lp/outputq
```

Don't remove the file, since it must exist with the right permissions and ownerships; truncating the file to zero length is safer. Now try `lpsched` again.

5. If the spooler locks up again, there is probably data corruption in the `pstatus` file; perform step 1 again. Now we need to make a list of all the printers currently in the system. Use `lpstat -v` to get this list, and make a permanent copy by redirecting it into a file as in

```
lpstat -v > /tmp/lp.list
```

If you are running Version 9.0 of HP-UX or have the spooler patch for 7.0 and 8.0, then you'll also have a list of printers that are remote or located on other systems.

The unprinted jobs in the spooler are located in the directory

```
/usr/spool/lp/request
```

where each printer has a directory of print jobs that have not been completed. Each job creates a pair of files: a short control file and the actual data for printing. If none of the print jobs are important, these files can be removed (details in the next step).

Otherwise, the individual data files can be moved to a temporary directory. Since many files may have been printed with special options, you should also move the control files to be used later in reprinting the jobs.

If `lpstat` did not provide a list of the remote printers with their remote computer hosts, you'll need to create a list of these printers manually.

Also note: If you have made customized model scripts in the

directory `/usr/spool/lp/interface`, make copies of these custom scripts in the directory `/usr/spool/lp/model` (the template directory). You may wish to rename the scripts in the interface directory since they will have the name of a printer.

Special note for JetDirect-connected printers and plotters: The model scripts for these printers are in a directory one level below the interface directory called *model.orig*. Be sure to save any customized scripts found in that directory too.

Also, the spoolkick procedure does not affect the bootp services used by JetDirect cards in this system; bootp is independent of the lp spooler.

6. With the spooler stopped, use `lpadmin -x` to remove every printer. If you have more than a few, you may wish to create a script to read the names of all the printers and run the `lpadmin -x` command automatically.

Once every printer has been removed, we need to move to the directory `/usr/spool/lp` to check on the remaining files and directories. Here is a summary of action steps:

- Remove the files *SCHEDLOCK* and *FIFO*.
- Clean the requests with `rm -r /usr/spool/lp/request/*`. Be sure to specify the `/*` at the end so that request will not be deleted.
- The following directories must exist but have **NO** files in them:
 - cinterface*
 - class*
 - info*
 - interface*
 - member*
 - receive*
 - request*
 - sinterface*
- The file *seqfile* can be removed or you can edit it to contain an ASCII number as the starting point for the next print request ID. (As a note, *seqfile* can be changed at other times but only when the spooler is shut down and there are no pending jobs to print.)
- The file: default should be zero-length.

- Zero out the status files with:

```
> pstatus (or you can use cat /dev/null > pstatus)
> qstatus (or cat /dev/null > qstatus)
> outputq (or cat /dev/null > outputq)
```

Do not remove these files. If they are removed, they must be recreated as zero-length files with the correct permissions and ownerships.

- Finally, we must verify that the existence, links, ownerships, and permissions are correct. With HP-UX 8.0 and higher, there is a very useful command: *pdfck*. With this command, you can verify the entire lp spooler file set. Just use

```
/usr/bin/pdfck /system/LP-SPOOL/pdf
```

This will report on any missing files or incorrect ownerships and permissions. Once *pdfck* finds no errors, we have completed the cleanup of the spooler. To verify that all is well, use:

```
/usr/bin/lpstat -t
```

The output should be something like:

```
/usr/bin/lpstat -t
scheduler is running
no system default destination
```

At this point, we can add the first printer, preferably one that is close by so we can verify operation immediately. Start by testing a directly connected printer (either serial or parallel). You can use SAM to add the printer or the commands

```
/usr/lib/lpshut
/usr/lib/lpadmin -p<prn> -v<devfile> -m<model>
/usr/lib/accept <prn>
/usr/bin/enable <prn>
```

where: *<prn>* is a name for the printer, *<devfile>* is the device file, and *<model>* is the model script to be used from the choices in `/usr/spool/lp/model`. Then, test the printer with

```
lp -d<printer-name> /etc/group
```


viNOT

THE UNIX TEXT EDITOR FOR **NORMAL** PEOPLE

- **Easy to use**
- **Easy to customize**
- **Easy on your budget**



For an evaluation copy or additional information, contact:



Computer Solutions, Inc.

120 E. Marks St. • Suite 225 • Orlando, FL 32803
407-649-0123 or 512-343-6634 • FAX 407-649-1407

All trademarks belong to their respective holders.

CIRCLE 110 ON READER SERVICE CARD

which will produce a one page (or so) listing of the group file.

This should completely restore operation of the lp spooler. Be sure that the latest lp spooler patches have been applied. For all the Op Systems:

You can obtain these patches from the HP Response Center, or electronically via HP Support Line. Connect to the Web service through <http://support.mayfield.hp.com>. ■

REV	PLATFORM	PATCH
7.0	300/400	PHCO_0545 (basic spooler fixes)
"	800	PHCO_0543 (basic spooler fixes)
8.0	300/400	PHCO_1433 (basic spooler fixes)
		PHCO_1968 (DTC printers)
"	800	PHCO_1432 (basic spooler fixes)
		PHNE_3526 (X.25 printing)
8.07	700	PHCO_1437 (basic spooler fixes)
		PHNE_3526 (X.25 printing)
9.0	300/400	(none)
9.x	700/800	PHCO_7033 (basic spooler fixes)
		PHCO_5898 (PCL5 scripts)
		PHSS_3223 (<i>hprpp</i> and <i>divpage</i>)
		PHCO_5370 (TSM printing)
		PHCO_4303 (<i>lpalt</i> and other problems)

Bill Hassell is an HP-UX systems support engineer in the HP Response Center in Atlanta, Georgia.

A close-up photograph of a wooden surface, possibly a desk or table. In the upper left, there is a small, white, classical-style column or pedestal. To its right, a metal object, possibly a hinge or a part of a mechanism, is visible. The lighting is warm and directional, creating strong shadows and highlights on the wood grain.

Sockets

Programming

by Andrew J. Phillips

This article provides an overview of programming with Berkeley Sockets on an HP-UX workstation. The level is technical but very basic, with lots of background and review material. More advanced readers may wish just to read the source listings. The concepts are applicable to almost any UNIX system, of course, and I'll provide some example programs written in C. I will discuss the terminology of networking and sockets, including TCP and UDP, and some client-server issues—including two different types of servers, and how sockets work.


Fundamentals and theory will be first, including terminology and other related topics that are needed to make appropriate decisions when programming with sockets. Then I'll cover the sockets API (Application Programming Interface), including various function calls and the data structures used by them. Once that has been accomplished, I'll present a simple as possible server to illustrate the various concepts with a concrete example. Next, I'll provide a real-life example of a

one-way UDP client-server system that is straightforward to program yet provides a service that is still in use today. The whole idea is to take the "magic" out of sockets programming and show that this sort of thing is not all that difficult, provided you understand the fundamental issues involved. If you're already sitting at a shell prompt, you can look at the files under `/usr/netdemo/socket` and `/usr/netdemo/af_unix/dgram` to see some example programs under HP-UX 9.X.

You should already understand programming concepts (UNIX environment is helpful) and also have some understanding of networks and operating systems. All examples will be given in C. I will make an effort to explain both the syntax and the source, but keep in mind there are always several different ways to toggle the same bits.

Networking with HP-UX

In this article, we will assume that networks can support



the DARPA Internet protocol suite, commonly referred to as TCP/IP. The IP portion of that moniker stands for *Internet Protocol*, which provides the packet delivery service (including routing) for other, higher level protocols such as UDP, TCP, and ICMP. At the IP layer, every packet can be considered independent of every other. This way, packets can be routed individually, since each contains both the source and destination address.

It is worthy of note that a single IP packet sent from machine A to machine B may arrive at machine B as several IP packets that, combined, contain the information in the packet sent from machine A. This phenomenon, referred to as *fragmentation*, can occur at the most inconvenient of times. It is also worthwhile to note that the IP protocol simply discards any packets that it finds in error (e.g., bad CRC checksum, too large, too small). It is up to the higher level protocols such as TCP, UDP, or ICMP, or the sockets programmer, to handle this situation, should it occur. I've used a lot of acronyms in this

paragraph, and it is important to understand what they mean and imply. So, the next section is a quick review of terminology, concepts, and protocols.

First, some definitions :

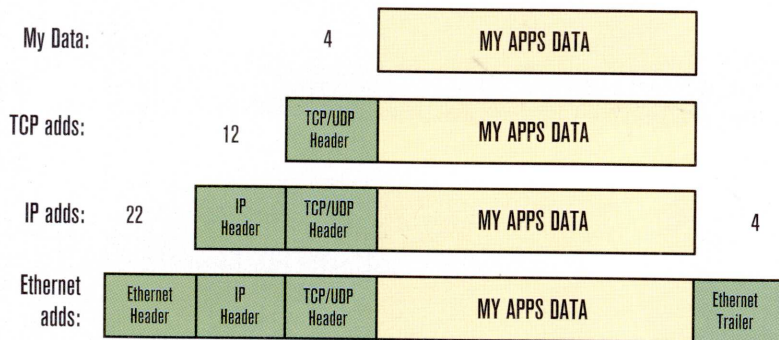
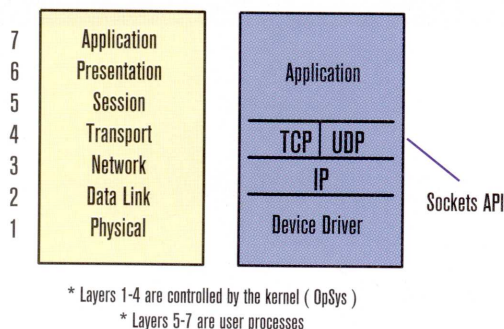
TCP is a connection-oriented protocol that provides a reliable, full-duplex connection between processes that is used like a "stream of bytes" from one process to another.

UDP is connectionless in that it provides only an unreliable message exchange medium for interprocess communication.

ICMP is yet another protocol used for providing status and control information. It is typically used for router configuration and *ping(1)*.

Terminology, Protocols, and Ports

Every TCP/IP implementation must support a minimum IP datagram size of 576 bytes. Many UDP applications (DNS,

FIGURE 1 Protocol Encapsulation**FIGURE 2** The Networking Picture

TFTP) only send up to 512 bytes of user data to stay within this limit and therefore guarantee that no fragmentation occurs, since this happens only when an IP datagram exceeds the underlying network's MTU (Maximum Transmission Unit—the largest “thing” that can be sent at one time). For Ethernet, this is about 5K or so. By definition, the maximum IP datagram size is 65,535 bytes, or 64K. The sockets programmer needs to be aware of these constraints when designing client-server systems so that performance will not be adversely affected.

Another concept to be aware of is the effect of protocol encapsulation, illustrated in *Figure 1*. This shows how much overhead is involved in sending a packet. In this example, there are no fewer than 42 bytes of overhead per packet. Since the minimum packet length for Ethernet is 72 bytes, there may even be additional “padding” of packets.

When dealing with networking protocols, it is sometimes helpful to refer to the ISO 7-layer model. See *Figure 2* for a depiction of how Berkeley Sockets

fits into this model. The model is meant to accommodate many protocols, and so our usage of it will have the three highest layers implemented in the application program. In fact, this is exactly the issue that will be addressed in the examples (i.e., session, presentation, and application).

A socket is an endpoint of communication. It is one side of a full-duplex connection. A connection is a communication link between two processes. An association is a 5-tuple (a group of exactly five things) that completely specifies the two processes that comprise a connection—for example, [protocol, l-address, l-process, f-address, f-process]. A half-association is also called a socket.

A connection is two sockets; each socket can both send and receive (full duplex in datacom parlance).

The port numbers and their classifications are illustrated in *Figure 3*. Port numbers below 1024 imply superuser, or root access, since these ports are designated as reserved. Ports are implemented by 16-bit positive integers in HP-UX. Some standard port numbers and protocols are 21/TCP for the ftp service and 23/TCP for the telnet service. Note that TCP and UDP port numbers are independent of each other. That is, there can be a UDP server on port 21/UDP and also a TCP server on port 21/TCP, and they both operate on different ports. Remember the definition of a socket? It included the protocol, so that is why the same port number can be used simultaneously with two different protocols. You can look at the file */etc/services* for a list of the protocols and ports already defined on your system.

Figure 3 also indicates that there is a group of what are known as “ephemeral” ports, numbered between 1024 and

Stuck using tar?

(or fbackup, cpio, or dump?)

What a sticky mess! Standard UNIX backup utilities force you to glue on scripts to make them work right, have performance like molasses, user interfaces that are clear as pitch, and reliability that could drag your career down into a pit. It's time to kick asphalt. Load BACKUP/9000 (it takes less than 10 minutes), and automate super fast, reliable backups and restores on any networked system via its slick user interface. Let BACKUP/9000 back up your Oracle databases hot, track tapes and files, manage media, schedule backups, etc. Don't get stuck - get something faster, easier, and more reliable. Contact us for a free demo today, before things get really messy.



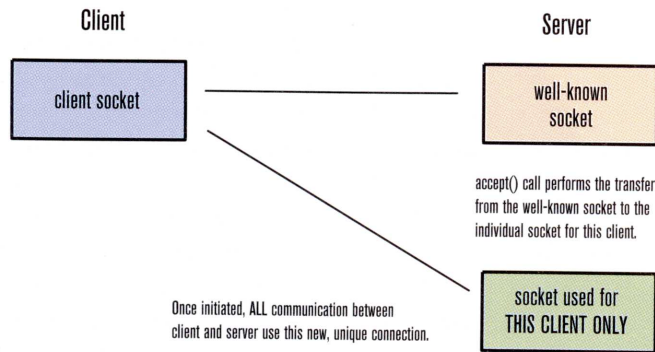
+1-800-89ORBIT
+1-510-837-4143
Fax +1-510-837-5752
info@orbitsw.com
www.orbitsw.com



FIGURE 3 HP-UX Port Numbers

- IANA = Internet Assigned Numbers Authority
- TCP ports and UDP ports are independent of each other, since the IP header specifies the protocol in use

1	1023	1024	5000	5001	65535
IANA Ports (reserved)		ephemeral ports		servers (nonprivileged)	

FIGURE 4 TCP accept() System Call

5000 inclusive. These ports are typically assigned by the system as the result of a special system call. This is explained in the next paragraph. The important concept to understand is that the particular number of an ephemeral port doesn't matter, because it will be communicated to the client process as part of the special system call. Ephemeral ports are used when a server will need to communicate with many clients simultaneously and will therefore need a socket for each client, in addition to the "well-known" socket that clients use to initiate a connection.

In order to implement a concurrent server, that is, a server that can provide service to more than one client at a time, the server must be able to handle several different ports at a time. Fortunately,

HP-UX (and UNIX in general) makes this relatively easy by providing the special system call *accept*.

The *accept()* call automatically allocates a random socket from a pool and associates this new socket with the server process. The server process then bases all future communications with the client process on this new socket. Using this method makes it much simpler to code a concurrent server, since the clients don't really care *which* socket is used by the server, as long as they can initiate the communication using the socket that they already know. An illustration of how this process works can be seen in Figure 4.

Client-server Basics

Client-server separates the process-

ing performed by an application program based upon real-world usage data. Thus, distributed programs using today's networks are more powerful than single processor systems. Client-server provides a flexible, scaleable, modular architecture for doing business. One of the problems with process-to-process communication is named the "rendezvous problem." This refers to the difficulty in synchronizing the actions of processes that not only execute at speeds measured in nanoseconds but are likely to be executing on different hardware as well. The rendezvous problem is solved by (1) having the server start first, and (2) having clients initiate connections. In this way, the server is usually "waiting" for clients, and the problem is reduced to one of connection establishment, rather than process synchronization. In order to do this, the client(s) must know where to find a server or servers before any type of connection can be made. So, we can define a *server* as a process that is waiting to be contacted by a client process so that the server can do something for the client.

Forms of Client-Server Communication

In general, client-server systems are of two basic types: connection-oriented and connectionless. The major dependency is whether the server needs to maintain its state between client requests. I will discuss both and give an example of a connectionless server in this article. I have also provided examples of several other more sophisticated servers and clients in the **Additional Code** section of the November *hp-ux/usr* on the Interex Web Site (<http://www.interex.org>).

In connection-oriented protocols, the clients send one or more requests over the connection to the server. This is

much like a file, or pipe between the two processes, providing the appearance of a “stream of bytes” between them. Connection-oriented protocols are best used when a request or response can be large, when the application protocol involves a continuing dialogue, or when the application requires reliable transport.

For systems that have one or more of these requirements, the transmission control protocol (TCP) is widely used. This protocol guarantees an error-free byte stream between the communicating processes. There is more overhead associated with using TCP, but many times the benefits of getting the correct data in the correct order outweigh the cost of the extra packets. Some of this extra cost is illustrated in *Figure 5*, which shows the packets involved in setting up and shutting down a TCP connection (the 3-way handshake). Note that these seven packets are required for every connection, regardless of how much data is sent over the connection. In general, systems using TCP have little or no control over the size or the number of packets transmitted.

For systems that are more message-oriented (where each message is independent of every other), a connectionless protocol can be used. This has the advantage of being lightweight (meaning minimal network traffic) and easier to program since the server can process each and every request independently.

The client places each request in an individual message and sends it to the server. The server then typically sends a response back to the client in a reply message. Be aware that this does *not* imply reliable transport, and only a “best-effort” message delivery service is pro-

vided. The protocol most often used for connectionless systems is UDP, also referred to as the “You Do it Protocol.” That means that the developer is responsible for any and all errors, including the following :

- a message never arrives at its destination
- messages are duplicated by the network
- messages arrive out of order

For obvious reasons, connectionless protocols are best suited for relatively simple systems. Some examples of systems that use UDP are SNMP (Simple Network Management Protocol) and DNS (Domain Name Service), although it’s worth mentioning that both of these systems can switch to a connection-oriented protocol when circumstances require (with large amounts of data or when security is a must).

Concurrency, Process Control, and Signals

For completeness, this section discusses some aspects of process control on UNIX systems. This is mainly required for discussing concurrent servers, but many client-server systems use a concurrent server, and it’s also good background material.

A new process is created only when an existing process executes the *fork()* function. *fork* creates a copy of the process that was executing. The process that called the *fork()* function is called the parent, and the new process is called the child. *Fork()* is the only function that is called once but returns twice (once in the parent, and once in the child), with different values. Next, a new program gets executed only when an existing process issues the *exec()* function. *exec* replaces the current process with the new program.

While most computers and operating systems today provide the appearance of concurrency (more than one program running at the same time),

FIGURE 5 TCP Connection Establishment and Shutdown





FIGURE 6 TCP Connection Sequence

true concurrency is achieved only with multiprocessor systems. UNIX provides apparent concurrency with time slicing. Server programs can use *fork* and *exec* to provide apparent concurrency for servicing more than one client at a time.

signal() is another widely used system call. This function allows a user-written event handler to be installed (and called) when the user process receives a signal. I have used the *alarm()* call along with the *signal()* call to provide “time-out” functionality for a UDP-based client server system. A code fragment that illustrates how to do this is shown below. First, we define a signal handler named *timeout()*, which first just re-installs itself and then increments a counter when called. Next, we define the signal handler in the main program. After that, we use the *alarm()* system call to set an alarm interrupt for 20 seconds in the future. Now, we can loop sending UDP packets to a server and waiting for a reply. If we send 5 packets without any reply, we will exit the loop and exit the program since the server is apparently unreachable.

```

/* This is the signal handler routine*/
static int timeout()
{
    signal(SIGALRM, &timeout);
    alarm( (unsigned long) 20 );
    packet_count++;
}

/* This shows how to use a signal handler. First we set up the
signal routine */
signal(SIGALRM, &timeout);
alarm( (unsigned long) 20 );

/* Now we enter the packet-sending loop */
while ( packet_count < 5 ) {

```

```

/* send UDP packets here */
if ( sendto() ) /* this means our send worked OK, and we can exit */
    return(0);
}

/* we only get here if all 5 packets timed out */
signal( SIGALRM, SIG_IGN );
return( -5 );

```

Using Berkeley Sockets-TCP

TCP sockets on most UNIX systems use the same conventions as files. That is, open ... read/write ... then close. This model works well because the socket is basically a “stream of bytes” just like normal files in the UNIX environment. The syntax for doing this is only slightly different from that for files, mostly because opening a socket requires information about the socket number, which host to connect to, and so on. There are a significant number of system calls related to sockets programming, so I have summarized some of them in *Table 1*. I will refer to these function calls throughout the rest of this article. Many of these will use data structures for passing data, so it is helpful to understand both the data structures and how the elements are used by the different functions. Some of these are specific to connectionless protocols (UDP) and some are for TCP protocols. I have used almost all of these in the examples, but I would suggest that the reader refer to the system man pages for more details about these functions on your system.

The sequence for setting up a TCP connection is illustrated in *Figure 6*. The arrows indicate where looping (or iterations) takes place. Calling sequences for both client and server are depicted, and the interaction is basically when one writes, the other reads, and vice versa. There is a little more overhead (in the form of more calls) on the server side because the server has to request a specific port number and then “listen” on that port for clients to connect.

In order to create a socket, one must initialize it first. This involves specifying such things as which protocol to use, which IP address to connect to, and which port number to use (the 3-tuple). The socket address structure is used to provide this functionality. A definition is shown below. The only complication here is that the IP address is also defined as a structure. So, the code defines two structures, an Internet address structure (*in_addr*), and a socket address structure (*sockaddr_in*). The definition for the socket address structure contains an Internet address.

TABLE 1

Berkeley Socket Calls

Socket	Create a socket & descriptor
connect	Connect to a remote peer
write	Send outgoing data
read	Get incoming data
close	Terminate all communication
bind	bind a local address to a socket
listen	place a socket in passive mode
accept	accept an incoming connection
recv	receive next incoming message
recvmsg	"
recvfrom	receive next message & get source address
send	send a message
sendmsg	"
sendto	send outgoing message to specified address
shutdown	terminate comm. in one or both directions
getpeername	obtain remote machines' address
getsockopt	obtain current options for a socket
setsockopt	change options for a socket

```

/* This is the Internet address structure */
struct in_addr {
    u_long    s_addr;          /* 32-bit netid/hostid */
};

/* This is the socket address structure */
struct sockaddr_in {
    short    sin_family;        /* AF_INET */
    u_short  sin_port;          /* 16-bit port number */
    struct in_addr sin_addr; /* 32-bit hostid/netid - Internet address */
    char    sin_zero[8];        /* unused */
};

```

The example below illustrates the use of the socket address structure. This example uses the *socket* call to create a socket and the *bind* call to bind the socket created to a port. Once the socket has been created (as a TCP socket in this case, using SOCK_STREAM), it then has various parameters set to (1) be an Internet socket, (2) accept connections from *any* other IP address, and (3) use a port number of 35000.

After this, the socket is ready for use.

```

/* include files for network stuff */
#include <netinet/in.h>
#include <sys/socket.h>

#define bzero(ptr, len)  memset((ptr), NULL, (len))

struct sockaddr_in myaddr;          /* declare a socket address
                                     structure */

int    mysocket;                    /* declare a socket descriptor */

mysocket = socket(AF_INET, SOCK_STREAM, 0); /* create socket */

bzero( &myaddr, sizeof(myaddr));    /* initialize the socket address
                                     structure */

myaddr.sin_family = AF_INET;          /* set up the address family to be
                                     Internet */

myaddr.sin_addr.s_addr = INADDR_ANY;  /* allow connections from ANYONE */

myaddr.sin_port = 35000;              /* set up the port number */

/* now, bind the socket */
bind(mysocket, (struct sockaddr *) &myaddr, sizeof(myaddr));

```


An Example Server

To provide the most basic example possible that is still a concrete one, the code shown below defines a “simple as possible” server using TCP. This server has no error checking whatsoever. First, there are several include files for defining data structures that will be used. There is also a definition for the port number that we will use, named *PORTNUM*. This server consists of a single *main()* function.

Within that *main* function, first two sockets are declared, one for the server and one for the client. Next, a data buffer is declared for receiving data from the client. Finally, a socket descriptor is declared. Now, we begin execution with the *socket* function call, which (as we already know) creates a socket. This socket is defined as being an Internet stream socket, which (to us) means the TCP protocol will be used. Next, we initialize the socket descriptor, specifying that the address family for this will be Internet, that we will accept *any* IP address that requests a connection, and finally that we will listen on TCP port number 9999.

Now we use the *bind* system call to bind this socket to a port (using our socket descriptor), and then we will listen for client requests. The 5 parameter in the listen call specifies that TCP will queue up to five clients requesting that port for us. If we get more clients than that before the server can service the first one, any further connections will be lost.

Next, we enter a do-while loop where we accept client connections on our socket, read 256 bytes from each client into our buffer, and then terminate the connection. We will do this forever, or until the server process is interrupted.

This illustrates the basic programming steps required to code an iterative TCP server.

```
/* NO error checking at all for this example, the idea is to make this as EASY to
read as possible. This is the "SIMPLE AS POSSIBLE" server */
```

```
#include <stdio.h>
#include <sys/types.h>
#include <sys/socket.h>
#include <netdb.h>
#include <netinet/in.h>
#define PORTNUM 9999
```

```
main ()
{
    int client_sd, /* this socket is for the client */
```

```
    server_sd; /* this socket is for the server's LISTEN socket */
    char buffer[256]; /* this declares a buffer to hold the client data */
    struct sockaddr_in name; /* this declares a socket descriptor */

    server_sd = socket(AF_INET, SOCK_STREAM, 0); /* this CREATES the socket */

    name.sin_family = AF_INET; /* these assignments will set up the descriptor */
    name.sin_addr.s_addr = INADDR_ANY;
    name.sin_port = PORTNUM;

    bind(server_sd, (struct sockaddr *) &name, sizeof(name)); /* this BINDS our socket */
    listen(server_sd, 5); /* put socket in passive (LISTEN) mode */
    do {
        client_sd = accept( server_sd, (struct sockaddr *) 0, (int *) 0);
        /* accept client's request for a connection */
        read(client_sd, buffer, 256); /* read the client's data into buffer */
        close(client_sd); /* terminate the connection */
    } while (1);
} /* end of main */
```

An Example One-Way UDP system

This entire client-server system is contained in *Listings 1-4*. These are (1) a common header file (*notify_tsl.h*), (2) server code (*collectd.c*), (3) client code (*notify_tsl.c*), and (4) an example data file.


This system was created to provide statistics about software usage. The client is basically a function that takes whatever information the user wishes to provide (in the form of two character strings), and sends it to the server (along with a local timestamp), where the information is written (along with a timestamp from the server and the client's IP address) to a log file. The log file can then be processed to see which IP addresses are using which applications, and when. Since this system is relatively simple and message-oriented, it uses UDP for communication from client to server. Note that this communication is one-way only, and the client performs a “blind send,” never knowing if its data was ever actually received. This was found to be acceptable in HP's network, since our error rate for a single UDP packet was .001-.002 percent. The server is iterative, since the time to process a client request is small (just write to the log file), and lightweight in terms of resource usage.

notify_tsl.h

This header file contains includes for various header files that are required to support different system features, along

Continued on Page 50

Today's the day
your overworked
processor spontaneously
combusts.



Are you prepared?

CLAM Associates can help you prevent the consequences of downtime. Power failures and surges cause 45.3% of data loss, but making your mission-critical applications continuously available is simple with CLAM Integration Services. Buying the work for you is *our* part. We solutions; our expertise in cluster reliable source to integrate into your system. Call (617) 621- or surf <http://www.clam.com> for quality services will provide



anything can happen.

Associates' MC/ServiceGuard software is the easy part; making it specialize in continuous availability architecture makes CLAM the most Hewlett-Packard's MC/ServiceGuard 2542, e-mail marketing@clam.com more information on CLAM. Our peace of mind in a world where

CLAM Associates. Solutions for a strange, strange world.

CIRCLE 6 ON READER SERVICE CARD

LISTING 1 notify_tsl.h

```

/*
   Filename: notify_tsl.h

This is the COMMON header file for the TSL STATS collection
client/server code modules, it contains GLOBAL defines, includes
and a TIMESTAMP function.
Written by Andy Phillips and Ricky Ralston 2/94.
*/
#include <stdio.h>
#include <stdlib.h>
#include <sys/socket.h>
#include <sys/utsname.h>
#include <sys/types.h>
#include <netinet/in.h>
#include <time.h>
#include <netdb.h>
#include <malloc.h>
#include <unistd.h>
#include <signal.h>

#ifndef SERVICE_NAME
#define SERVICE_NAME      "collect" /* service in /etc/services */
#endif

#ifndef COLLECTOR_PORT
#define COLLECTOR_PORT      22375 /* UDP port number for service */
#endif

#ifndef COLLECTOR_SERVER_NAME
#define COLLECTOR_SERVER_NAME "example.atl.hp.com" /* STATS server */
#endif

#ifndef DAEMON_HELLO
#define DAEMON_HELLO      "TSL STATS Collector ver B.02 daemon started "
#endif

#ifndef LOGFILE
#define LOGFILE            "/usr/appstats/tsl.app.data"
#endif

#ifndef BUFFERSIZE
#define BUFFERSIZE        512 /* maximum size of packet */
#endif

#ifndef SERVER_DATA_SEPARATOR
#define SERVER_DATA_SEPARATOR " " /* for writing LOG file */
#endif

#ifndef CLIENT_DATA_SEPARATOR
#define CLIENT_DATA_SEPARATOR "@" /* for WS-COE data */
#endif

#ifndef TIMESTAMP_FORMAT
#define TIMESTAMP_FORMAT    "%Y-%m-%d %T %a %Z " /* shared between C/S */
#endif

#ifndef TIMESTAMP_LENGTH
#define TIMESTAMP_LENGTH    30
#endif

extern int errno;

void timestamp ( loctime )
char *loctime;
{
    time_t timer;
    struct tm *timeptr;

    /* get local timestamp */
    time(&timer);
    timeptr = localtime(&timer);
    strftime(loctime,TIMESTAMP_LENGTH,TIMESTAMP_FORMAT,timeptr);
} /* timestamp */

```

LISTING 2 collectd.c

```

/*
   COLLECTD.C

This program implements the collectd server for TSL stats
collection. It listens on udp port 22375 ( which is in
etc/services ) for datagrams from clients. Each datagram is
processed iteratively by extracting the source IP address
from the datagram along with the data, writing this to a
log file ( /usr/appstats/tsl.app.data ), then appending
a timestamp.
*/

#include "notify_tsl.h"

/*
   M A I N

This routine starts the server. It forks, leaving the child
to do all the work, so it does not have to be run in the
background. It sets up the socket, and for each incoming
packet, it writes a record. Each packet consists of an
application name, version number, and associated data.
The server will extract the source IP address from each packet
and write this, along with the applications name,version, and data
to the logfile. A timestamp is then appended to the record, and the
transaction is complete.
*/

int packet_cnt ;
char begin_time[TIMESTAMP_LENGTH];

void show_status()
{
    char nowtime[30];
    char msgbuff[255];

    (void) signal(15,show_status);
    timestamp(nowtime);

    sprintf(msgbuff,"Collectd started on %s \nENDING on %s \n, %ld packets.\n",
        begin_time,nowtime,packet_cnt);

    printf("%s",msgbuff);
} /* show_status */

main(argc, argv)
int argc;
char *argv[];
{
    int addrlen;
    char *inet_ntoa();
    char *startup_time;
    char *ptr;
    char *ProgName;

    char loctime[TIMESTAMP_LENGTH]; /* holds timestamp */
    int i; /* loop counter for signals */
    int s; /* socket descriptor */
    int cc; /* contains the number of bytes read */
    char buffer[BUFFERSIZE]; /* buffer for packets to be read into */

    struct hostent *hp; /* ptr to host info 4 requested host */
    struct servent *sp; /* pointer to service information */

    struct sockaddr_in myaddr_in; /* for local socket address */
    struct sockaddr_in clientaddr_in; /* for client socket address */
    char *reqaddr; /* for client host's address */
    FILE *fd, *fopen(); /* for the LOGFILE */

    /* Get program basename */
    if((ptr = strrchr(argv[0], '/')) == NULL) {
        ProgName = argv[0];
    } else ProgName = ++ptr;

```


LISTING 2 collectd.c, continued

```

/* init counter */
packet_cnt = 0;

/* Ignore all signals */
for(i = 1; i <= NSIG; i++) signal(i, SIG_IGN);

signal(15, show_status);

/* clear out address structures */
memset((char *)&myaddr_in, 0, sizeof(struct sockaddr_in));
memset((char *)&clientaddr_in, 0, sizeof(struct sockaddr_in));

/* Set up address structure for the socket. */
myaddr_in.sin_family = AF_INET;

/* The server should receive on the wildcard address,
 * rather than its own internet address. This is
 * generally good practice for servers, because on
 * systems which are connected to more than one
 * network at once will be able to have one server
 * listening on all networks at once. Even when the
 * host is connected to only one network, this is good
 * practice, because it makes the server program more
 * portable.
 */
myaddr_in.sin_addr.s_addr = INADDR_ANY;

/* Find the information for the "collect" server
 * in order to get the needed port number.
 */
if ((sp = getservbyname (SERVICE_NAME, "udp")) == NULL) {
    fprintf(stderr, "%s: %s not found in /etc/services\n",
        ProgName, SERVICE_NAME);
    exit(1);
}

/* Verify the port information */
if (sp->s_port != COLLECTOR_PORT) {
    fprintf(stderr, "%s: %s should be on port %d not %d\n",
        ProgName, SERVICE_NAME, COLLECTOR_PORT, sp->s_port);
    exit(1);
}

/* Now, set the servers' port address */
myaddr_in.sin_port = sp->s_port;

/* Create the socket. */
if ((s = socket (AF_INET, SOCK_DGRAM, 0)) == -1) {
    perror(ProgName);
    fprintf(stderr, "%s: unable to create socket\n", ProgName);
    exit(1);
}

/* Bind the server's address to the socket. */
if (bind(s, &myaddr_in, sizeof(struct sockaddr_in)) == -1) {
    perror(ProgName);
    fprintf(stderr, "%s: unable to bind address\n", ProgName);
    exit(1);
}

/* Now, all the initialization of the server is
 * complete, and any user errors will have already
 * been detected. Now we can fork the daemon and
 * return to the user. We need to do a setpgp
 * so that the daemon will no longer be associated
 * with the user's control terminal. This is done
 * before the fork, so that the child will not be
 * a process group leader. Otherwise, if the child
 * were to open a terminal, it would become associated
 * with that terminal as its control terminal. It is
 * always best for the parent to do the setpgp.
 */
setpgp();

switch (fork()) {
case -1: /* Unable to fork, for some reason. */
    perror(ProgName);
    fprintf(stderr, "%s: unable to fork daemon\n", ProgName);
    exit(1);

```

LISTING 2 collectd.c, continued

```

case 0: /* The child process (daemon) comes here. */
    /* Close stdin, stdout, and stderr so that they will
     * not be kept open. From now on, the daemon will
     * only report to the LOG file. This daemon
     * will loop forever, waiting for requests and
     * responding to them.
     */

    /* Open log file */
    if ((fd = fopen(LOGFILE, "a")) == (FILE *)NULL) {
        perror(ProgName);
        fprintf(stderr, "%s: Unable to open %s.\n", ProgName,
            LOGFILE);
        exit(1);
    }

    /* Show server daemon is OK upon invocation */
    timestamp(loctime);
    strcpy(begin_time, loctime); /* save startup time */
    startup_time = malloc(strlen(DAEMON_HELLO)+sizeof(loctime));
    sprintf(startup_time, "%s %s", DAEMON_HELLO, loctime);
    fprintf(stdout, "%s", startup_time);
    close(stdin);
    close(stdout);
    close(stderr);

    fprintf(fd, "%s\n", startup_time);
    fflush(fd);
    free(startup_time);

    /* The SERVER code is what happens inside this loop */
    for(;;) {
        /* Note that addrlen is passed as a pointer
         * so that the recvfrom call can return the
         * size of the returned address. We don't
         * need this info. anymore, but it's good
         * to understand anyway.
         */
        addrlen = sizeof(struct sockaddr_in);

        /* This call will block until a new
         * request arrives. Then, it will
         * return the address of the client,
         * and a buffer containing its request.
         * BUFFERSIZE - 1 bytes are read so that
         * room is left at the end of the buffer
         * for a null character.
         */
        if ((cc = recvfrom(s, buffer, BUFFERSIZE - 1, 0,
            &clientaddr_in, &addrlen)) == -1) exit(1);

        /* Make sure the message received is
         * null terminated, the recvfrom returns
         * the length of the data received.
         */
        buffer[cc] = '\0';

        /* get local timestamp */
        timestamp(loctime);

        /* Extract the source IP address from the
         * datagram.
         */
        reqaddr = inet_ntoa(clientaddr_in.sin_addr);

        /* Now, write the time, IP address and the data to the
         * logfile record and flush to disc */
        fprintf(fd, "%s %s %s\n", loctime, reqaddr, buffer);
        fflush(fd);
        packet_cnt++; /* count packets processed */

    } /* for loop */

default: /* Parent process comes here. */
    exit(0);
} /* case stmt */
} /* end collectd */

```


LISTING 3 notify_tsl.c

```

/*
 *
 *      NOTIFY_TSL.C
 *
 *      This program uses a datagram socket to transfer 2 char strings
 *      to the TSL server r3125sv2.atl.hp.com on port 22375. It sends a single
 *      UDP packet (< 512 bytes) and does NO error checking. It can be used
 *      either as a C function call ( this code ), or as a "wrapper"
 *      that can be called from a shell script or the command line.
 *
 */

#include "notify_tsl.h"

int notify_tsl ( app, appdata )
char *app;
char *appdata;
{
    int s;                /* socket descriptor */
    char loctime[TIMESTAMP_LENGTH]; /* holds timestamp */
    struct hostent *hp;    /* pointer to host info for server host */
    struct sockaddr_in myaddr_in; /* for local socket address */
    struct sockaddr_in servaddr_in; /* for server socket address */
    static char stats_data[BUFFERSIZE]; /* buffer for counter data */

    /* clear out address structures */
    memset ((char *)&myaddr_in, 0, sizeof(struct sockaddr_in));
    memset ((char *)&servaddr_in, 0, sizeof(struct sockaddr_in));

    /* Set up the server address family to be Internet. */
    servaddr_in.sin_family = AF_INET;

    /* Get the host information for the server's hostname
     * assume that DNS (or /etc/hosts) is setup and works OK */
    hp = gethostbyname (COLLECTOR_SERVER_NAME);
    if (hp == NULL) return(-1);

    /* Set the IP address using the data structure from gethostbyname */
    servaddr_in.sin_addr.s_addr = ((struct in_addr *) (hp->h_addr))->s_addr;

    /* Set the port number. */
    servaddr_in.sin_port = COLLECTOR_PORT;

    /* Create the socket. */
    s = socket (AF_INET, SOCK_DGRAM, 0);
    if (s == -1) return(-2);

    /* Bind socket to some local address, a port
     * number of zero will be used so that the system will
     * assign the first available port number. An address
     * of INADDR_ANY will be used so we do not have to
     * look up the internet address of the local host.
     */
    myaddr_in.sin_family = AF_INET;
    myaddr_in.sin_port = 0;
    myaddr_in.sin_addr.s_addr = INADDR_ANY;
    if (bind(s, &myaddr_in, sizeof(struct sockaddr_in)) == -1)
        return(-3);

    /* get local timestamp */
    timestamp(loctime);

    /* init DATA to be sent timestamp+name.version+data */
    sprintf(stats_data, "%s %s %s", loctime, app, appdata);

    /* Send the request to the server. */
    if (sendto (s, stats_data, strlen(stats_data), 0, &servaddr_in,
                sizeof(struct sockaddr_in)) == -1)
        return(-4);
    return(0);
} /* end notify_tsl */

```

LISTING 4 Example Data File

```

TSL Collector daemon started on : Mon Dec 13 14:47:00 EST 1993

1993-12-13 14:50:32 Mon EST 192.41.7.106 1993-12-13 14:48:35 Mon EST ANDY1STTEST.00
here is my DATA

1993-07-15 07:46:23 Mon EDT 192.50.117.143 1993-07-15 07:47:02 Mon EDT CDB.MOD.D.05
Runs: 1 Adds: 0

1994-01-14 09:32:30 Fri EST 192.41.7.109 1994-01-14 09:32:08 Fri EST NSC Admin Tool
(LEEMAH) USER: leemah_client

1994-01-14 14:30:11 Fri EST 192.41.7.108 1994-01-14 14:31:25 Fri EST fview.alpha
Uses: 1 Reads: 1 Help: 1

1993-03-15 03:18:03 Fri EST 192.163.192.44 1993-03-15 09:18:12 Fri MET PostnoteXm:
4.03 Machine: 9000/710, HP-UX: A.09.0

1996-03-15 06:16:09 Fri EST 192.50.133.187 1996-03-15 06:16:24 Fri EST Helpdesk
Status Board rev. 4.11 MACHINE: i3125g17

1996-03-15 06:48:32 Fri EST 192.32.121.44 1996-03-15 06:48:46 Fri TSL CMSMGR A.02.04
ZAM CMSDB

1996-03-15 06:57:45 Fri EST 192.32.121.44 1996-03-15 06:57:59 Fri TSL CMSMGR A.02.04
MD CMSDB

1996-03-15 07:54:04 Fri EST 192.50.113.43 1996-03-15 07:55:20 Fri TSL SMLOGON
A.01.01 ISE094:SYS.MANAGER

1996-03-15 07:55:52 Fri EST 192.50.113.43 1996-03-15 07:57:08 Fri TSL SMLOGON
A.01.01 ISE117:SYS.MANAGER

```

with global defines that are used by both client and server (to help keep them in synch). I strongly suggest using a common header file when developing client-server systems. It aids greatly in keeping the application protocol the same between client and server. Lastly, a timestamp function that is used by both client and server is defined.

collectd.c

This is the server code. It includes the common header file, declares a counter and a buffer, and defines an interrupt service routine. The rest of the code is a single main function that provides all of the server functionality. Within this function, variables are declared for sockets, counters, and buffers, as well as a log file for saving data from each transaction. Next, variables are initialized, and interrupts are ignored except for signal 15, which writes the current time and packet count to the log file.

Next, we use *getservbyname()* to check */etc/services* for the configured port number. This allows us to change the server

port by editing the `/etc/services` file. Now we can create and initialize the server's socket and bind it to a port. After this, the server is ready to become a daemon, and so it will perform several systems calls to facilitate this. Once the child process has successfully become a daemon, it opens the log file, writes a startup message, and begins an infinite for loop. This loop contains a `recvfrom()` call to get both that data and the IP address from clients. If the `recvfrom` call fails, the server terminates; otherwise, it will just wait indefinitely for clients to send messages. Every time `recvfrom()` is successful, a timestamp is generated, the IP address of the client is extracted from the UDP packet, and the client data along with this information is written to the log file. The packet count is then incremented, and the loop continues.

notify_tsl.c

This contains the client code. Basically, this function (which takes two character string parameters) simply puts a timestamp and the two parameters into a UDP packet and sends it.

It was designed to be as lightweight as possible (so end users would not be affected). The code contains variables for the socket and two address structures, along with a local buffer. First the hostname is looked up (so we can use DNS), the socket is created and bound, and a timestamp is generated. Once that has been accomplished, a data buffer is loaded with the client data and the timestamp, and then this buffer is sent in a UDP packet to the server.

Example data generated by this system can be seen in *Listing 4*. Note the two timestamps (one from the client and one from the server) that allow

Protect Your Data & Pick Up the Pace with ...

RAID RUNNER

Don't fall behind the pack. Safeguard your computer's data with RAIDRUNNER, the disk array that clears RAID price-to-performance hurdles with ease. Get the safety & performance of RAID 3 and 5, plus capacities that go the distance – from 5 GB to over 30 GB.

**What are you waiting for?
On your mark, get set, go...
RAIDRUNNER!**

Supported Operating Systems

- HP-UX • MPE • Digital UNIX®
- Windows NT™ • Windows • NetWare®
- AIX • SunOS • Solaris • and others

*The Pacesetter in RAID Protection & Value
for Networking and Desktop Computing*

U.S.A. toll-free, call **800.237.4641**
Outside the U.S.A. dial **408.364.6500**
or e-mail us at **info@bering.com**

©1996 Bering Technology – All rights reserved. All brand names and product names are trademarks or registered trademarks of their respective holders.

CIRCLE 4 ON READER SERVICE CARD

detection of time zone differences (and clock skew).

This concludes this article on sockets programming. If you have further questions regarding this material, you may e-mail me. My address is ajp@atl.hp.com. ■

Andrew J. Phillips is an information technology engineer within the AFO Technology Solutions Lab in Atlanta, Georgia. In the 11 years he has worked for HP, he has been a telecommunications engineer, a system manager, and a program analyst.

REFERENCES

- [1] *UNIX Network Programming*, W. Richard Stevens, Prentice Hall, Englewood Cliffs, New Jersey, 1990
- [2] *Internetworking with TCP/IP—volume II*, Douglas E. Comer and David L. Stevens, Prentice Hall, Englewood Cliffs, New Jersey, 1991
- [3] *UNIX Networking*, Stephen G. Kochan and Patrick H. Wood, Hayden Books, Carmel, Indiana, 1989
- [4] *HP C/HP-UX Reference Manual*, HP Part No. 92453-90024, Third Edition, 1992
- [5] *Berkeley IPC Programmer's Guide*, HP Part No. 98194-90012, First Edition, 1989
- [6] "Client-Server Development," Antoinette Agracewicz, *UNIX Review*, June 1994, vol. 12 no. 6
- [7] *The design of the UNIX Operating System*, Maurice J. Bach, Prentice Hall, Englewood Cliffs, New Jersey, 1986
- [8] *TCP/IP Network Administration*, Craig Hunt, O'Reilly and Associates, Sebastopol, CA, 1992

Book Review

by Chris Curtin

The Dilbert Principle

EASILY THE MOST IMPORTANT book about managing a company since *Reengineering the Business*. A no-holds barred, cut to the chase, "we're all in this together" look at corporate America. This book is written by one of the most influential authors on management in our time, so you'll want to keep a copy of it on your desk to know how to evaluate, understand, and respond properly to the memos and edicts from senior management, which strives to make your jobs easier and more enjoyable.

Yeah. Right.

For those who don't know who Scott Adams or Dilbert are, pick up this book at your local bookstore, open to page 61 and read the first couple of paragraphs of "Your Input Is Important to Us." Then look at the cartoon on the bottom of page 188. Or get your daily paper and look for the *Dilbert* comic strip. It is usually in the Business section. Or look at the geeky, engineering types at the local mall or technical bookstore. At least 100 will have a Dilbert T-shirt on. (I can call them "geeky, engineering types" because I am one of them. ;-)

The table of contents pretty much describes the tone and approach of the book:

Forward: Big Opening

Introduction: Why Is Business So Absurd?

Chapter 1 The Dilbert Principle

Chapter 2 Humiliation

Chapter 3 Business Communication

Chapter 4 Great Lies Of Management

Chapter 5 Machiavellian Methods

Chapter 6 Employee Strategies

Chapter 7 Performance Reviews

Chapter 8 Pretending to Work

Chapter 9 Swearing: The Key to Success for Women

Chapter 10 How to Get Your Way

Chapter 11 Marketing and Communications

Chapter 12 Management Consultants

Chapter 13 Business Plans

Chapter 14 Engineers, Scientists, Programmers, and Other Odd People

Chapter 15 Change

Chapter 16 Budgeting

Chapter 17 Sales

Chapter 18 Meetings

Chapter 19 Projects

Chapter 20 ISO 9000

Chapter 21 Downsizing

Chapter 22 How To Tell If Your Company Is Doomed

Chapter 23 Reengineering

Chapter 24 Team-building Exercises

Chapter 25 Leaders

Chapter 26 New Company Model: OA5

Scott Adams was an engineer, living in a cubicle, when he started drawing Dilbert, an engineer, living in a cubicle, working for an idiot boss. (Scott insists that most of his material is from e-mail he has received. Do you believe him?)

Unlike the other Dilbert books, this book is not just a collection of old comic strips you can read all at the same time instead of having to wait for tomorrow's newspaper (or Web site update). In this book, Adams explains (or attempts to explain) many of the actions to which management subjects its employees.

The commentary is sprinkled with humorous footnotes, comics from the daily series, and actual e-mail from readers, minus the headers and employee and company names to protect the stupid. I very much enjoyed his style of writing, often insulting himself (check

At-a-Glance

Title:

The Dilbert Principle: A Cubicle's-eye View of Bosses, Meetings, Management Fads, & Other Workplace Afflictions

Author:

Scott Adams

Publisher:

United Feature Syndicate, Inc.,
1996, 336 pages. ISBN: 0-88730-787-6

Price:

\$20.00

out the footnote on page 197) and leaving nothing sacred.

While humor is the delivery vehicle, a lot of what he is presenting is unbelievable. Even looking at the topics he covers and assuming some degree of exaggeration, the policies and tasks of management he describes made me cringe several times.

In the last chapter, the author changes direction completely and describes what he sees as the perfect company. After reading about all the horrible, stupid, and dangerous actions companies take, he opens himself up for criticism by defining a model he would use to run a company. Not to spoil the ending, do *not* read the last chapter first. It has a greater impact when read after the rest of the book.

Okay, maybe I'll ruin the ending a little bit. The best idea presented in the whole book is: (Guess what, I can't print it here. After all this is a "family" magazine. See page 321, Item #1.)

As for the opening paragraph of this review, I'd think about your management's style before leaving this book on your desk. They may realize you are on to them. ■

Chris Curtin, a software developer for Bradley Ward Systems, Inc. in Atlanta, Georgia, specializes in device driver development for factory automation on the HP 9000. He can be reached via e-mail at: chris@bwilab3.atl.ga.us.

Don't Let Poor System Performance Drive You Crazy!



LUND
PERFORMANCE SOLUTIONS

Performance Beyond Expectation
240 2nd Ave. SW, Albany, OR 97321, USA

System managers today are being faced with complex problems, performance management is one of them. **SOS/9000 Performance Advisor** is the ultimate on-line performance tool for HP 9000 systems.

With instant access to performance data as it happens, system managers can optimize performance, identify bottlenecks, increase throughput and solve performance problems quickly and easily.

For more information or a **FREE 30 day trial copy**, call today at: (541) 926-3800!




SOS/9000
Performance Advisor

CIRCLE 129 ON READER SERVICE CARD

ALPHANUMERIC PAGING FOR UNIX

**ROBUST, RELIABLE,
USER-FRIENDLY DELIVERY
OF MESSAGES ANYTIME,
ANYWHERE**

- Email forwarded to pager automatically
- Pages can be generated from scripts and network monitoring programs
- GUI and command line interface
- Works with any paging service
- Automatic email confirmation, history logs and error reporting
- Client-server technology
- Works with digital and alphanumeric pagers



**Personal Productivity Tools
for the Unix Desktop**

14141 Miranda Rd
Los Altos Hills, CA 94022
Email: sales@ppt.com
Tel: (415) 917-7000
Fax: (415) 917-7010
<http://www.ppt.com>

CIRCLE 91 ON READER SERVICE CARD

by *Bob Combs*

Installing NT 4.0

THE INSTALLATION OF A new operating system can be fairly intimidating if you haven't gone through it before. But like most tasks, it's really quite straightforward once you've done it. There are a few things you need to check and a few things you should know before starting. I'll detail what to expect to get your Microsoft Windows NT 4.0 installed.

Before You Begin

Before you begin, you should refer to Microsoft's Hardware Compatibility List (HCL). The HCL manual is included in the box with Windows NT, is online at Microsoft's Web site, and is on the Microsoft TechNet CD-ROM. It lists all the known computer models and peripheral devices supported by Windows NT. I always use this manual to look up my computer, disk, CD-ROM, video card, and SCSI controller (if any), before starting any installation. Some of the devices have footnotes associated with them. Pay particular attention to these, since they usually point out potential problems with that particular card or device. If you have a disk, controller, CD-ROM, or video board that is not listed on the HCL, the odds are you may not be able to generate a system on that computer.

The computer should have a minimum set of resources. Microsoft claims that the machine should have a minimum of 12 MB of RAM and 75 MB of free disk space. Microsoft recommends, however, that the computer have 16 MB of RAM. I've found that a full system will use about 125 MB of disk space and needs at least 16 MB of RAM. I prefer 32 MB of RAM for my workstation, which improves the speed of most programs. If you're planning on installing Microsoft Office and a development language, you'll want at least a 1-GB disk drive.

If you've checked your hardware on the HCL and verified that you have enough resources, you'll need to note a few items before you start the installation.

- | | |
|------------------------------|---|
| ■ Computer name | a name you assign to your computer |
| ■ Administrator password | your local super-user password |
| ■ Network IP address | the IP address assigned to your computer |
| ■ Network Subnet mask | the subnet mask associated with your IP address |
| ■ Default gateway IP address | the IP address of your segment's router |
| ■ DNS IP address | the IP address of the DNS you use |

The first two items, computer name and administrator password, are needed for all installations. The last four are needed only if your computer is connected to a network that uses TCP/IP. The last two are needed if your network is connected to the Internet. If you have any questions about the last four, ask your server administrator or local network services person.

Ways to Install

Windows NT can be installed from Windows 3.1, from another version of Windows NT, or onto a new blank system. While you can install Windows NT from floppies, I'm assuming a CD-ROM for installation.

Installing to a blank new system means you'll use the three (3) setup floppies that come with the CD-ROM. (You can always recreate these floppies from the CD-ROM

using "WINNT /O".) Boot the Setup floppy and follow the directions to install NT.

Installing on top of Windows 3.1 or Windows NT can be performed without the floppies by running the installation program directly from the CD-ROM. If the CD-ROM drive is drive letter D:, and we're installing onto an Intel-based computer, then the `D:\i386` directory contains the installation files. You can install Windows NT onto an existing Windows 3.1 or Windows for Workgroups 3.11 system (16-bit systems) by running `D:\i386\WINNT`. If you are installing onto an existing Windows NT or Windows 95 system (32-bit systems), run `D:\i386\WINNT32`.

The standard operation of WINNT is to (1) create setup floppies, (2) copy files to a temporary disk directory (`WIN_NT.~LS`), and then (3) boot the setup floppy to start the installation. Since we already have the setup floppies, and would prefer not to have to load them anyway, I use the parameter switch `WINNT /B` to run without the floppies, i.e., don't create them and don't load them.

Towards the end of the installation process, you are asked for a blank floppy to create an "Emergency Repair Disk." This floppy is very important for repairing the configuration and accounts of your particular system. But because you'll be adding additional items to the system after the initial installation, you'll probably have to recreate the repair disk again later. Therefore, I tend to skip this process during the initial installation by adding the `/X` switch to WINNT. My typical install string looks like `D:\i386\WINNT32 /B /X` to start the install process from Windows 95 or Windows NT. Later I create the repair disk by running the NT program `RDISK` from the Start/Run prompt.

Installation Process

Starting WINNT/WINNT32 from the CD-ROM will copy files to the hard disk and then ask to restart the computer.

After restart, a special limited installation NT system boots up. It detects hardware and chooses drivers for the devices it detects. You must choose from options such as Upgrade, Repair, and New Installation, depending upon whether a previous NT system is detected on the disk. I want to point out that when it detects mass storage devices, ensure that your CD-ROM or interface is listed. If not, type 'S' to specify additional devices and add it to the list. Otherwise the install will fail a few steps later when it looks for floppies instead of the CD-ROM.

Also, if you have a multiprocessor computer, then when the Computer Configuration is presented, highlight the "computer" definition and press enter. You'll then be asked for the manufacturer's driver disk, which contains the multiprocessor HAL driver supplied with your machine.

You'll be asked to format a partition or accept the current partition to install into. The default answers will guide you if you're unsure. When asked the question of which file system to use, FAT or NTFS, I answer NTFS. The FAT file system is compatible with MS-DOS, Windows 95, or OS/2, allowing dual boot configurations. I prefer NTFS because it allows full NT security.

The install program then scans the hard drive for corruption and then restarts the computer again.

Windows NT 4.0 has a new look when the computer reboots this second time. It has a new installation wizard to make installation easier using the Next/Back dialog boxes. After you

accept the license terms, files are copied to the NT directories.

When asked for the installation type, choose "Typical" unless you know what you're doing.

You'll be asked for your name, company name, and then the name of your computer. Then you'll be asked for the password you're assigning to the Administrator account. Don't forget this password!

You have a choice of creating the Emergency Repair disk or not, and then are asked about networking. Choose the networking connection types (modem or direct connect) that you'll use; and the protocols (TCP/IP, IPX, or NetBEUI) in use on your network.

If you are asked about DHCP, answer "no" unless you know that your IP addresses are dynamically served and support DHCP.

If you use modem dial-up networking, answer "yes" to RAS networking. You'll then be asked to detect the modem and configure it.

Finally, you are asked to select your time zone, and the system cleans up. The Emergency Repair disk is then created, if selected, and the system restarts to bring up your Windows NT system.

Pressing Ctrl-Alt-Delete will pop up the logon dialog box, where the user name of Administrator and the password you selected can be entered. You'll notice that the 4.0 shell looks like Windows 95. Have fun exploring the similarities and differences between Windows 95 and Windows NT 4.0! ■

Bob Combs is the Director of Systems Architecture at PCSI in Englewood, New Jersey, a company specializing in client-server technology. He is a Microsoft Certified Systems Engineer (MCSE) and holds a master's degree in electrical engineering.

by Joseph Berry

Software Activities Picking Up

A NUMBER OF INTERESTING pieces of software were announced on Usenet during the last two months. Activities are picking up. My Web site (such as it is) has changed. You can now find me described at <http://www.topshot.com/joeberry.html>.

COMP.INTERNET.NET-HAPPENINGS

weblint (v 1.017)

This program proves that you need to keep your finger on many Usenet newsgroups or you will miss some nice programs. Weblint is a Perl program that analyzes html files for syntax errors. A wonderful tool for Web page developers, weblint is similar to the lint syntax checker for C programs. Weblint requires Perl 4 or 5 to run. This software, by Neil Bowers <neilb@cre.canon.co.uk>, is available via the Web at <http://www.cre.canon.co.uk/~neilb/weblint/> or via ftp as <ftp://ftp.cre.canon.co.uk/pub/weblint/weblint-1.017.tar.gz>. It installs in a few seconds.

GNU.EMACS.SOURCES

e-mh-alias.el

I do my e-mail using the mh package (see the Goodies column in the March 1996 issue). While I sometimes use mh directly, I more typically use it within the editor emacs. This emacs elisp file extends the basic mh offering by doing automatic name completions for aliases in the mh-letter buffer. This is for lazy people such as myself.

rmime.el

This is another great little mh add-on for emacs users. More and more often I get mime-encoded e-mail. These are e-mail messages that have self-

describing descriptors added to the beginning, allowing you to send pictures and other kinds of non-textual documentation. With rmime, you simply enter control-c twice on the line that describes the enclosure and the appropriate viewer is executed showing you what was sent.

COMP.SYS.HP.HPUX

tcpdump v 3.2

In the previous issue of *hp-ux/usr*, I discussed the program tcpdump (Version 3.0) and mentioned that although this is a very useful and powerful tool, it did not run on HP-UX systems. This latest version of tcpdump and its companion libraries libpcap (Version 0.2) and bpf (Version 1.1) from the Network Research Group at the Lawrence Berkeley National Laboratory now support HP-UX. The three packages are available via ftp from the following site: <ftp://ee.lbl.gov>. The files are in the root directory and are called *tcpdump.tar.Z*, *libpcap.tar.Z*, and *bpf.tar.Z*.

tom v 0.01

This sounds like an interesting program. It is the first release of an object-oriented programming language, inspired by Objective-C that doesn't (so it is claimed) suffer a C heritage. The program comes to us from Pieter Schoenmakers <tiggrr@es.ele.tue.nl>.

This first release has been tested on HP-UX systems; it also runs on Linux (X86) and FreeBSD (X86) systems. Information about tom is available at <http://tom.es.ele.tue.nl:8080/>.

A few of the features of tom listed below are based on the announcement of the program in the Usenet newsgroup:


Usability. With tom, the semantics of

SUPRTOOL IS HERE!

HP 3000's favorite database utility is now on HP-UX.

With Suprtool you can

- ▶ select and sort Oracle data.
- ▶ export data to MS Access, and to your choice of report writers.
- ▶ provide easy migration for HP 3000 data.



Call Robelle today to find out how Suprtool can increase database efficiency on your HP 9000.

Suprtool, the Data Handyman for HP-UX

Robelle
CONSULTING LTD.

Unit 201
15399-102A Ave.
Surrey, British Columbia
Canada V3R 7K1

Toll-free: 1-800-561-8311
Phone: (604) 582-1700
Fax: (604) 582-1799
E-mail: info@robelle.com
WWW: <http://www.robelle.com>



Official Robelle Distributors

Australia, New Zealand.....61 2 484 3979
Central America.....502 2 314786
Czech, Slovak Republics.....42 2 723305
France, Belgium.....33 1 69 86 60 00
Germany.....49 7621 689 190

Greece, Italy.....30 1 777 0561
Holland, Belgium.....31 13 521 56 55
Hong Kong.....852 2609 1338
Mexico.....52 5 813 1325
Saudi Arabia, U.A.E.....966 1 477 4555

Scandinavia.....46 8 683 00 50
Singapore.....65 441 2688
South Africa.....27 21 685 7809
Switzerland, Austria.....41 31 981 06 66
Taiwan.....886 2 545 2166
United Kingdom, Ireland.....44 171 473 2558

Qedit is a trademark of Robelle Consulting Ltd.

CIRCLE 100 ON READER SERVICE CARD

the language do not change with the availability of sources. Put more strongly, the possibilities of the language are not restricted in any way if sources are not available.

This in contrast to, for example, Eiffel, where the availability of only the short form of a class severely hampers the possibilities of subclassing.

As another example, given a C++ library, suppose you want to replace a class by a different one. In that case, the member variables must remain identical if referred to from within one of the other objects in the library. The same is true for virtual methods if invoked from the other objects. If any of the member functions of the class to be replaced is in-lined in the other objects, the possibility of successfully completing this venture within the original goals is further reduced.

Dynamic. At runtime, tom provides full typing information on objects, their instance and class variables, and arguments to and return value from methods. All methods are dynamically bound. Methods are provided to read and set instance and class variables given their name, to retrieve an object's class, to query a class for its position in the inheritance hierarchy, to send an object a message of which the selector need not be known at compile time (a bit like function pointers in C++, but being dynamically bound much more flexible—more like perform in SmallTalk and Objective-C, but without restrictions on the types and number of arguments).

Storage Management. All objects are allocated from the garbage collected heap. This implies that it is impossible to have problems with dangling references, previously valid pointers suddenly pointing into the void, memory leaks, etc.

Return Values. In most languages,

there is an arbitrary limit on the number of values a method or function can return without reverting to tricks. This limit is 1, and pointer arguments are needed to return more than one value, or the return values need to be packed into an aggregate entity which can then be returned as the single return value. In Common Lisp, multiple-valued returns are possible, but they are special.

Tom imposes no limit on the number of values that can be returned from a method.

Exceptions. In most languages (C++, Java, NeXT's Objective-C libraries, to name but a few), an exception handler is executed in its own context instead of the context of the exception raiser, implying that it is impossible for the handler to direct the actions of the condition raiser—the stack has already been unwound and nothing can be done about that.

The tom exception mechanism is modeled after Common Lisp conditions. A tom condition can be signaled; if none of the condition handlers is interested in handling a condition, execution continues in the context of the condition signaler, as if the signal were a normal method invocation. A condition handler can direct the condition signaler by setting the value to be returned from the signal invocation. It can also decide to perform a non-local exit, the equivalent of a C long jump.

COMP.SYS.SUN.APPS

SciTeXt

I know this package exists since I've looked at the Web page associated with it. I have not tried the software, however, so I cannot attest to its quality. Nevertheless, SciTeXt is a new, free word processor (as opposed to text processor) for UNIX systems.

Some of the more impressive features include the following: WYSIWYG, support for PostScript fonts, paragraph and document templates, auto-correction and spell checking, footnotes, table of contents, indices, layouts for tables, charts, graphics, etc., etc. The software comes with a Tcl/Tk help browser as well as a Tcl/Tk installation script to make the installation process go easy.

The package is freely available to anyone for commercial or non-commercial use, and can be ftp'ed from the following site:

<http://www.uni-paderborn.de/~SciTeXt/archives/getSciTeXt.html>

You can also go to the home page where a description of the software is available (<http://www.uni-paderborn.de/~SciTeXt/>).

Unfortunately, no precompiled version for HP-UX is available. The sources are downloadable and based on what I read, building an HP-UX version shouldn't be a problem (it does require X11R6, however). Binaries for IRIX, Linux 2.0, and Solaris 2.5 are currently available.

COMP.LANG.PERL.ANNOUNCE

PGPPerl

This is a Perl module that contains a suite of PGP security routines. Although the library is quite usable, it is still a long way from being complete, and the interface is subject to change.

This module contains the following functions: PGP conventional encryption, PGP public key encryption/decryption, and PGP key generation.

The module can be downloaded from <http://www.systemics.com/software/>.

COMP.UNIX.MISC

Amanda v 2.3.0

The Advanced Maryland Automatic Network Disk Archiver (amanda) package comes courtesy of the University of Maryland at College Park, where it is used daily at the Computer Science Department. This package, built on top of standard backup software (using GNU tar, gzip, etc.), is a backup system designed to archive many computers on a network to a single large-capacity tape drive.

The software includes tape management and automatic backup checks; the results can be e-mailed and installation and configuration are easy. In addition, amanda supports Kerberos authentication software and optional automatic data compression.

Amanda can be acquired by anonymous ftp from <ftp.cs.umd.edu>. Go to directory `/pub/amanda` and look for the latest copy.

WEB Pages

<http://developer.netscape.com/library/one>

This site presents two white papers, a technical overview of the Netscape Open

Network Environment (ONE) application architecture and a strategy paper on the company's intranet product line. It's interesting reading and presents Netscape's vision of the future of Web software.

<http://www.byte.com>

Byte Magazine's Web site offers a really nice benchmarking tool that you might want to download and look at. I have found it to be very useful. In addition, many resources, including articles and interviews are made available to the Web surfer.

<http://Web.mit.edu/asante/www/siteday.html>

Those of you who have checked out my Web site know that my background included studying astronomy as a university undergraduate. I have recently come back to that subject as an amateur and am finding it immensely pleasing to see so many Web sites on the Internet dedicated to the field. This site presents a different Web site every day that is dedicated to the field of astronomy. ■

Joseph Berry is a senior software developer at Landmark Systems Corporation in Vienna, Virginia. He is one of the authors of Landmark's PerformanceWorks products, PerformanceWorks/Smart Agents for UNIX. A former HP 3000 systems specialist for Hewlett-Packard, he has been in the computer industry for more than 25 years. He can be reached at joe@wayne.unix.landmark.co.



HP WORLD '97

The Conference & Expo
for Hewlett-Packard
Computing Solutions



Keeping You A Step Ahead

AUGUST 24 - 29, 1997 • CHICAGO, ILLINOIS • NAVY PIER

HP WORLD '97 Conference & Expo will focus on the day-to-day challenges faced by users and provide unique "close-ups" of the most compelling technologies. This gives you an even greater opportunity to keep a step ahead of your computing challenges. You'll find solutions for:

Technology Close-ups

- ▼ HP-UX
- ▼ MPE/iX or MPE/V
- ▼ Real Time
- ▼ Windows NT
- ▼ Development Tools
- ▼ Networking
- ▼ PC
- ▼ Databases
- ▼ Emerging Technologies
- ▼ Internet Technologies
- ▼ OpenView

Today's Challenges

- ▼ Enterprise-wide data access
- ▼ Electronic commerce
- ▼ Multi-platform integration/development
- ▼ Security/high availability/system management
- ▼ Management issues
- ▼ Desktop integration

Save the Date!

Mark your '97 calendar today! Or to learn more, check out our web site at <http://www.hpworld.org> or call 1-800-INTEREX (outside the US: 1-408-747-0227).



by Larry Headlund

Last X Release

THE LAST RELEASE OF X will hit the streets at the end of 1996, God willing and if the creek don't rise. The X Consortium, descendent of the original X development team at MIT, will disband, and responsibility for X Window will be transferred to the Open Group, a subsidiary of the Open Software Group, the people who brought you Motif. A lot of the X development people are going over to the Open Group at the same time. I am writing this in August of 1996 and developments by the time you read this may make me look foolish, but that has never stopped me before.

What Is Going On?

The release of X numbered X11R7 will be the last release of X. This release will include the technology known as Broadway, about which there will be more information later. From other developments it seems that X, Motif, and CDE (Common Desktop Environment) will all be under one roof. This last has rung a lot of people's chimes.

The X Consortium is being absorbed by the Open Group, not the other way around. When the announcement was first made, some asked if it meant that Motif would now be available under the same terms that X had been, full source and free redistribution. No way. When this fact was absorbed, it was immediately followed by a cry that free X was being taken away.

This is a more complicated question. In the discussion that followed (on *comp.windows.x* primarily), Bob Scheifler, the original author of X and president of the X Consortium, pointed out that X had used the free distribution and source code as a means to wide distribution and use of the technology, not as an

end in itself. Contrast this with the position of the Free Software Foundation, where the notion of software being free is central, nay overriding.

For those just tuning in, the Free Software Foundation, the GNU (GNU not UNIX) project, the GNU General Public License, and the GNU General Public Library License are creations of Richard Stallman and dedicated to the proposition that software should be free in the sense that the source code should be available to any user and that any user should be able to make changes and redistribute those changes. Note that nothing here is said about monetary issues, although this principle has economic consequences. "Free" is used in the sense of *libras*, not in the sense of *gratis*, clearing things up for you classical scholars.

The GNU Project is the creation of an entire UNIX-like operating system that is free in this sense. The goal was first announced, along with the principles behind it, in 1984 in Stallman's GNU manifesto. Discussion of the whole project's wisdom, implementation, ramifications, and consequences hasn't stopped yet.

The GNU General Public License is a tool to ensure the goal is reached. Briefly, if you receive software covered under the GPL, you have the right also to have the source code. You have the right to redistribute both without any further permission from the author of the software. You have the right to modify the software and distribute that too; however, if you redistribute the software (and what is meant by "redistribute" is a frequent subject for debate), your modifications are also covered by the GPL. You can't take GPL software, add your own code, and redistribute while keep-

ing your changes proprietary. All the software is covered by the GPL. All means all. You couldn't write an interface to a database for a GPL spreadsheet, for example, unless the interface code and the database code was available.

This restriction is in effect even if you yourself don't have access to the offending code. You couldn't write and distribute an interface to a third-party database unless the third party also released its code. I keep saying write and distribute since what you do on your own computer is unrestricted. Only when the result is distributed does the license come into effect. This is a rather strict interpretation of the GPL. Some authors would permit the distribution I have just described. The reason for this restriction is to block end runs around the GPL. If it wasn't there, you could have company A with proprietary code, and company A.1 making use of GPL code and Company A's libraries to produce software that was unusable without paying tolls to company A.

A similar license, the General Public Library License, isn't so stringent. With this license, you can link and use GPLL code with proprietary code, provided you allow access to the original library. This relaxation was a consequence of the use of the enormously popular gcc C compiler. Under a strict interpretation of the GPL, any code compiled with gcc and linked with the gcc libraries would be covered by the GPL. This dampened use of the gcc considerably. This particular feature of the GPL resulted in the (erroneous) belief that just using GNU tools resulted in your code becoming "infected." The gcc libraries are covered under the General Public Library license so this is not true for gcc. It is also not true for the other

Consider the advantages of the **INTEREX DIRECT MAIL SERVICE**

- ◆ 46,000 Hewlett-Packard computer users in our database
- ◆ A one-week turnaround time
- ◆ Less than 2% returns because of incorrect addresses
- ◆ 74% of our users are their company's decision makers
- ◆ Full refund for all nixes (bad addresses)
- ◆ 15% discount on all orders over 10,000 labels
- ◆ List updated weekly

We offer the most value for your money through direct mailings and the INTEREX Card Deck.

Contact Phil Nguyen, INTEREX,
P.O. Box 3439, Sunnyvale,
CA 94088-3439, USA

(800) INTEREX • (408) 743.4619 FAX; (408) 747.0947

tools such as bash, the gnu shell, or any other GNU tools. If you write your code using the Emacs editor, the Free Software Foundation has no claims on you! There is one interesting exception. The exception is bison, the GNU version of yacc. Since yacc and the rest of the herd function by inserting part of themselves into the output, code run through bison is covered by the GPL.

There is a class of exceptions to the proprietary library linkage question, even for strict GPL programs. These are parts of the "standard operating environment." I put that in quotes because what qualifies is not based on principle, but on what is usually part of (UNIX) operating systems. These include compilers and libraries, shells, and windowing systems. You see why I had the UNIX in parentheses: compilers are not a normal part of PC operating systems but were for UNIX systems when GNU was conceived. You can link GNU software to software which exploits the quirks of PC compilers. But it is the windowing system exception that brings us back to X.

X has become the standard for GNU software. GNU partisans are passionate about their software and their cause. Note that X is not covered by the GPL. There is nothing to keep a company from taking X code, making modifications, and keeping the results proprietary. It is done all the time. HP starts with the code from the X consortium, makes modifications, hopefully enhancements, and keeps the results for its own use (and ours). Just because it is based on X code doesn't mean we, as users, have a right to it. Contrast this with releasing X under the GPL. The source code from HP as well as the X-based Motif code would be available and freely redistributable by all. The big

question is, Would this now proprietary code have been written at all under those conditions? Those who want to participate in the never-ending debate on these questions should check out *gnu.misc.discuss*. Bring your asbestos suit.

This relationship between X and GNU, and the fact that X has been free in both senses of the word since its birth fueled some of the anxiety about X being changed and particularly about the relationship with the OSF. Motif has not been universally popular with free software partisans. Primarily, the reason is that its source code is not free. There has been a tendency for GNU partisans to prefer Athena widgets. By no means all GNU developers—there are many excellent Motif-based GNU programs, but they have been avoided by the purists. This is mixed with a general anti-OSF feeling. The roots of this are more complicated.

Remember that the OSF produced Motif through a union of Sun Microsystems' competitors, HP prominent among them; Motif was counterpoint to Sun's Open Look. Sun was very popular with many GNU partisans, not least because they had the same roots, Berkeley Software Distribution (BSD). Sun's OS was BSD-based while HP was System V with BSD enhancements, IBM's AIX and DEC's UNIX were System V, etc. Ironically, the OSF was formed in response to agreements between Sun and AT&T that resulted in Sun moving to Solaris, which was System V.4 based. The BSD crowd had a natural sympathy for Sun in these lamented interface wars. Sun solidified its position by releasing its Remote Procedure Call (rpc) tools in the public domain. Free software partisans tend to regard the Open Software Foundation with suspicion.

What has all this to do with the price of bread? As Bob Scheifler pointed out, X11R5 is still available under its original distribution terms—do anything you want, just acknowledge the contribution—and X11R7 will be distributed under the same terms. If the free software community wants to pick up the ball, they are free to do so. It is the X consortium's position that the underlying X technology is now mature enough that it doesn't need an organization dedicated to extending the technology. Just as X11R6 moved the Athena widget set out of the core distribution and into the contrib section, in effect all of X has been moved to the contrib section.

One advantage of X, Motif, and CDE being under one roof is hopefully more coordination. Since CDE was introduced, we have had an X release (X11R6), a Motif release (Motif 2.0 based on X11R5), and the public release of CDE (based on X11R5 and Motif 1.2). Is there something wrong with this picture? What am I supposed to use? With one organization responsible for all three products, we should get the inconsistencies taken care of.

But maybe I am an optimist. ■

Larry Headlund is president of Mathematical Engineering, Inc. and has been involved with commercial UNIX since 1982 and with X since 1988. He can be reached at lnh@world.std.com or (617) 242-7741.

InterWorks '97

Conference & Exhibition

- System Administration Conference
- Software Development Conference
- Internet / Intranet Training
- Windows NT Highlights

APRIL 12 – 17, 1997
PHILADELPHIA, PA

<http://www.InterWorks.org>

CIRCLE 63 ON READER SERVICE CARD



CSL Perspective

AFTER WRITING MY LAST column regarding job skills, I realized that there is a huge number of possible topics, many of which cannot be covered in a single column or short discussion. I do want to continue that discussion, but first I have some very exciting information to pass on to our members.

At the HP WORLD conference in August, Interex and the InterWorks Technical Users Forum formally announced the FAST (Freely Available Software Technology) project. FAST consists of a group of dedicated Interex members who will be porting the top 100+ freely available software packages to HP-UX 10.0. The packages will be organized into bundles to make it very easy to find and load the software that you need.

With the release of 10.0 by HP in 1995, we have had many inquiries from members about particular packages and the availability on the new HP-UX operating system. Although the Interex CSL cannot guarantee any particular package will run on any particular operating system version, our experience indicates that most will. Although we believe that our packages would run under 10.0, we believe that this is a great opportunity to improve our depth and breadth of programs and to make great improvements in the configuration, packaging, and bundling of the library. We also want to take advantage of some of the enhancements in the operating system such as the System V.4 file layout standards and the new Software Distributor.

Many of you have been regular users of the HP Public Domain Archive at Liverpool. This organization has done some really fantastic work in making a wide variety of packages available to the HP-UX user community. If you look at

their catalogue, you will find over 1,100 packages, all of them built and tested under HP-UX 9.x. Imagine what kind of effort it might take to move that volume of software onto 10.0. It clearly will require a broad-based effort, with a much larger team of people working together. This team of porters, working from a published distribution standard (you can view it at <http://www.cs.berkeley.edu/~staelin/software.html>) will start from a list of the 100+ most popular packages in such categories as image-handling, mail/news, networking tools, software development, system administration, word processing, and World Wide Web. These packages will then be available to the HP-UX users through various sites.

For more information on the FAST effort, consult the project Web page at <http://www.interworks.org/library/FAST.html>.

The FAST project effort plays into the job skills topic I would like to discuss—Teamwork. In July, I participated as an adult advisor on a backpacking trip to Philmont Scout Ranch in New Mexico. You notice, I was an advisor, not a leader. Each crew of scouts consists of 8 to 10 boys with up to four adult advisors. This group of people forms a team with a single scout selected as the crew leader. This leader must be capable of making decisions, achieving consensus, and resolving differences. Some say that's a pretty tall order for a 16-year-old, and you would be right.

Fortunately, it is the crew, operating as a team, that really empowers that leader and also empowers the team to be more than the just the sum of its parts. As our trek started to wind down and individuals began to falter, we could fall back on each other to get through those tough days. We could rely on the strength of the team to pull us through.

More and more these days, our day-to-day work activities require us to contribute to as well as lead teams. We may use other terms like "group," "committee," "task force," or "working group" to describe collections of individuals brought together to accomplish some goal, but none of these terms seems to capture all the nuances of "team." A team is more than just a collection of individuals; it also has as part of its makeup the collection of skills, insights, and experiences of those individuals. It is that blending of capabilities which brings the team its potential for success.

As teamwork becomes more important to our success as system professionals, it is also becoming a skill to look for in people joining our organizations. Unfortunately, teamwork skills are not on the top 10 list of college taught skills. Teamwork is one of those soft skills best developed through trial-and-error experience. Some of the qualities of a good team member include:

1. Good listening
2. Clear, concise communication
3. Ability to make and keep commitments
4. Ability to defend personal preferences and ideas
5. Willingness to compromise your preferences for the good of the group
6. Willingness to assist others when required
7. Being prepared to follow the team leader, or assume leadership when needed

Among the strategies that can be used to assess these qualities, I have found role playing to be very effective. In role playing you propose a situation

<http://www.iodatasys.com>

Email: iodata@interramp.com

Total Internet

- Statistical Reporting
- Site Hosting
- On-Line Ordering
- Fax-Back Info

Faxing Solutions

- Host-Fax - HP 3000
- Replix - HP 9000/UNIX
- LanFax - PC Networks

DataComm Consolidation

- Bridges, Mux, CSU/DSU in one expandable box

High Availability

- Disk Mirroring
- Disk Shadowing

I/O Data Systems, Inc.



HP 3000 HARDWARE

HP 9000 HARDWARE

HP CERTIFIED



I/O DATA SYSTEMS, INC.

Ph. (216) 835-2211 Fax (216) 835-0220

"An Inc. 500 Company"

CIRCLE 171 ON READER SERVICE CARD

to the individual and then ask how he would deal with it. Use some team experiences of your own to build up these scenarios, putting the person in another's shoes. Also assess leadership skills by outlining a conflict between team members in the scenario. Glean from the individuals some of their ideas and preferences and then challenge them with a scenario in which they would need to compromise. By assessing the person's experiences on teams, and then asking them specific questions about their role, you can quickly determine possible areas for improvement or areas of particular strength. Don't be afraid to share your assessment with the person later on.

Teamwork is a powerful concept. It can be a springboard to success or a formula for disaster. When you are in the back country for 10 days, you have little

opportunity for assessment; you need to make it work. Don't get caught with a team that can't pull together; do your homework and you'll be amazed at what will happen. ■

Paul Gerwitz is chairman of the CSL committee and is a technology specialist at Eastman Kodak Company in Rochester NY. He can be reached at (716) 477-3067 or e-mail at gerwitz@interex.org or gerwitz@kodak.com.



HP 1000 Guru

Q: I just discovered the existence of the PASSWORD command in GRUMP.

Actually, one of my users discovered this for me, and in the process noticed he was allowed to change the password for any user on the system. Of course, he still needed to know the user's password to be able to change it, but still, this behavior doesn't seem right to me. Only superusers should be able even to attempt a PASSWORD command on other users. The PASSWORD command for normal users should be restricted to allow them only to change their own password, not the password of other users.

A: Yes, this is true. GRUMP gives normal users the capability to change the password for any user on the system.

The workaround for this is easy if you are using SECURITY/1000. It turns out that if SECURITY is being used, then a user must have a capability of 25 (still not super(31)) to use the PASSWORD command on other users. If a capability 20 user attempts the password command, the following error is reported:

```
GRUMP> PA WALT
```

```
Capability level not high enough to perform PASSWORD command.
```

So as long as users have a capability of less than 25, they will not be able to use the PASSWORD command. Otherwise, see the next question.

Q: Can I change the capability level required to use the PASSWORD command in GRUMP, and how is this done?

A: Buried in the *System Manager* manual lies the complex answer to this question. Modifying capability levels requires the following:

1. EDIT the file: `/vcplus/security.tbl`

This file contains all the default capability levels for system calls and commands. For the GRUMP commands, look around line 432 for the following:

```
passwd 0 25 0 0 *change user password (1)
```

Change the value 25 to 31. This is the required capability.

2. Run STGEN on this modified file as follows:

```
CI> STGEN SECURITY.TBL - - -
```

This will create the files:

```
SECURITY.LST SECURITY.MAC SECURITY.REL
```


3. Regenerate your system, using this new *SECURITY.REL* in the System Messages section of your answer file.
4. Reboot using this new system file.

This process can be used to modify any existing capabilities in *SECURITY/1000*.

Q: Is it possible to create an RTE readable tape using *dd* on an HP 9000?

A: Yes. Let's assume you have an FST archive file on an HP 9000 system and need to move this to an HP 1000 that has no network access.

You can use *dd* on the 9000 to create a tape that will be readable by FST on the 1000. Use *dd* as follows:

```
$ dd if=filename of=/dev/rmt/tapedrive obs=10240
```

where:

filename is the FST archive file on the HP 9000

tapedrive is the device file for the tape drive

(Note: the DDS tape drives must be similar on both machines. If the 9000 tape drive uses compression, the 1000 tape drive must be capable of reading compressed tapes.)

In this example, *obs=10240* sets the output file block size to 10,240 bytes. This only affects the performance, especially for dds tape drives. It is not required.

Once *dd* has written the tape, it can be read using FST on the HP 1000 just like any normal FST backup tape.

Q: Is there a 'tail' utility for the 1000?

A: New in the CSL this year is a tail program that emulates the HP-UX version. Or you can use the following simple script:

```
*
* The following script runs LI on the specified file name
* in $1, and displays the last 'window' (20 lines).
*
li -> /$1/ $1
*
```

LISTING 1 Prog1

```
$CDS ON$
$HEAP 2$
PROGRAM prog1(input,output);

label 99;

const
    alloc_flags = 0;
    dealloc_flags = octal('40000'); { unlock }

type
    int = -32768..32767;
    btype = packed array[1..16] of char;
    rec_t = record
        val : int;
        str : packed array[1..10] of char;
    end;
    ptr_rec = ^rec_t;
    flag_rec = record
        word1 : int;
        word2 : int;
    end;

var
    p1 : ptr_rec;
    p2 : ptr_rec;
    err : int;
    flags : flag_rec;

function xema1_alloc(flags : flag_rec; var err : int)
    : ptr_rec; external;
function xema2_alloc(flags : flag_rec; var err : int)
    : ptr_rec; external;
function RteReturnShema(addr : ptr_rec; flags : int) : int; external;

procedure exec_7 $alias 'EXEC'$ (x:int); external;

begin
    flags.word1 := alloc_flags;
    flags.word2 := 0;
    p1 := xema1_alloc(flags, err);
    if p1 = nil then
        begin
            writeln('prog1: init #1 failed, err=',err);
            goto 99;
        end;
    p2 := xema2_alloc(flags, err);
    if p2 = nil then
        begin
            writeln('prog2: init #2 failed, err=',err);
            goto 99;
        end;

    p1^.val := 11111;
    p1^.str := 'aaaaaaaaaa';
    writeln('val #1 =', p1^.val);
    writeln('str #1 =', p1^.str);

    p2^.val := 22222;
    p2^.str := 'bbbbbbbbbb';
    writeln('val #2 =', p2^.val);
    writeln('str #2 =', p2^.str);
    writeln('wait for prog #2 to finish, then type GO');

    exec_7(7);

    writeln('val #1 =', p1^.val);
    writeln('str #1 =', p1^.str);
    writeln('val #2 =', p2^.val);
    writeln('str #2 =', p2^.str);

    flags.word1 := dealloc_flags;
    err := RteReturnShema(p1, flags.word1);
    if err <> 0 then
        writeln('prog1: dealloc #1 failed, err=',err);

    err := RteReturnShema(p2, flags.word1);
    if err <> 0 then
        writeln('prog1: dealloc #2 failed, err=',err);

99:
end.
```


LISTING 2 *Prog2*

```

$CDS ON$
SHEAP 2$
PROGRAM prog2(input,output);

label 99;

const
  alloc_flags = 0;
  dealloc_flags = octal('40000'); { unlock }

type
  int   = -32768..32767;
  btype = packed array [1..16] of char;
  rec = record
    val : int;
    str : packed array [1..10] of char;
  end;
  ptr_rec = ^rec;
  flag_rec = record
    word1 : int;
    word2 : int;
  end;
var
  ptr1 : ptr_rec;
  ptr2 : ptr_rec;
  flags : flag_rec;
  err : int;

function xema1_alloc(flags : flag_rec; var err : int)
  : ptr_rec; external;
function xema2_alloc(flags : flag_rec; var err : int)
  : ptr_rec; external;
function RteReturnSchema(addr : ptr_rec; flags : int) : int; external;

procedure exec_7 $alias 'EXEC'$ (x:int); external;

begin
  writeln('wait for prog #1 to suspend, then type GO');
  exec_7(7);

  flags.word1 := alloc_flags;
  flags.word2 := 0;
  err := 0;
  ptr1 := xema1_alloc(flags, err);
  if ptr1 = nil then
    begin
      wrteln('prog2: alloc #1 failed, err=', err);
      goto 99;
    end;

  ptr2 := xema2_alloc(flags, err);
  if ptr2 = nil then
    begin
      writeln('prog2: alloc #2 failed, err=', err);
      goto 99;
    end;

  writeln('prog2: val #1 = ', ptr1^.val);
  writeln('prog2: str #1 = ', ptr1^.str);
  writeln('prog2: val #2 = ', ptr2^.val);
  writeln('prog2: str #2 = ', ptr2^.str);

  ptr1^.val := 99;
  ptr1^.str := 'ccccccccc';
  ptr2^.val := 88;
  ptr2^.str := 'ddddddddd';

  writeln('prog2: val #1 = ', ptr1^.val);
  writeln('prog2: str #1 = ', ptr1^.str);
  writeln('prog2: val #2 = ', ptr2^.val);
  writeln('prog2: str #2 = ', ptr2^.str);

  flags.word1 := dealloc_flags;
  err := RteReturnSchema(ptr1, flags.word1);
  if err <> 0 then
    writeln('prog2: dealloc #1 failed, err=', err);
  err := RteReturnSchema(ptr2, flags.word1);
  if err <> 0 then
    writeln('prog2: dealloc #2 failed, err=', err);

  writeln('prog2 done!!');
99:
end.

```

LISTING 3 *FORTTRAN Subroutine*

```

ftn7x
$cds on

integer*4 function eaddr(a)
integer*4 a
eaddr=a
return
end

Sema /xema1/
c xema1_alloc will allocate an extended schema of 1 page
integer*4 function xema1_alloc(flags, err)
implicit none

integer*4 flags
integer*2 err

integer*2 e1(1024,1) ! Space for sec schema #1, 1024words*1 page
common /xema1/ e1

integer*2 size
integer*2 RteAllocSchema
integer*4 eaddr

size = 1 ! We will be sharing 1 page
err=RteAllocSchema('xema1name', size, e1, flags)
if(err.ne.0) then
  xema1_alloc=0
else
  xema1_alloc=eaddr(e1)
endif
end

Sema /xema2/
c xema2_alloc allocates another schema
integer*4 function xema2_alloc(flags, err)
implicit none

integer*4 flags
integer*2 err

integer*2 e2(1024,1) ! Space for sec schema #2, 1024words*1 page
common /xema2/ e2

integer*2 size
integer*2 RteAllocSchema
integer*4 eaddr

size=1
err=RteAllocSchema('xema2name', size, e2, flags)
if(err.ne.0) then
  xema2_alloc=0
else
  xema2_alloc=eaddr(e2)
endif
end

```


HP 3000 • TRAINING • NEW • USED • FULL WARRANTY

HP 9000

LEADER

RS 6000

X TERMINALS
PLOTTERS
PRINTERS
DAT • UPGRADES
OEM • 700'S
DISK • 800'S

BUY • SELL • RENT • REPAIR • PARTS

1 800 553 0592

FAX 612-476-1903

wdpibmhp@mr.net
Established 1987

WORLD DATA PRODUCTS

7400 Flying Cloud Drive • Eden Prairie • MN • 55344

CIRCLE 128 ON READER SERVICE CARD

Q: We have a legacy HP 1000 application written in Pascal. We would like to utilize the new for 6.0 Multiple Shared EMA capability, but are having a hard time getting it to work from Pascal. Are there any examples of this specifically for Pascal?

A: Yes. The HP 1000 lab has come through with sample programs for Pascal that use the new Shema model. The PASCAL programs require a FORTRAN subroutine.

There are two Pascal programs that communicate through two shemas. Each contains FORTRAN routines that declare the common blocks and call RteAllocShema. Prog1 initializes the data, then calls exec(7) to wait for the second program to reset it. Prog2 calls exec(7) to wait for the first program to set the data. These are in Listings 1-4. ■

Walt Boeninger works in the HP Response Center in Mountain View, California. He has been supporting the HP 1000 for 15 years. His e-mail address is: walt@hpwrcxe.mayfield.hp.com.

LISTING 4 Link Command Files

LINK command file for PROG1

```
sz,32
em,,1
es,xema1,2
es,xema2,3
li,pascal_cds.lib
re,ppas1.rel
re,pftn.rel
en,prog1
```

LINK command file for PROG2

```
sz,32
em,,1
es,xema1,2
es,xema2,3
li,pascal_cds.lib
re,ppas2.rel
re,pftn.rel
en,prog2
```




by Anita Harris

Questions and Answers

Q: I am writing an application for HP-RT 2.11 and 2.20 that uses POSIX functions in C++. I find that when I compile and link my program using:

```
$HPRTroot/hpux/bin/CCrt prog.C -oprog -lp
```

I get undefined symbols that should have been defined in the POSIX library *libp.a*. *libp.a* should have been searched by specifying *-lp* in the compile string. What is wrong?

A: The problem you are running into can be resolved by using the following compile/link string:

```
$HPRTroot/hpux/bin/CCrt prog.C -oprog -lC -lp
```

This will ensure that the HP-RT version of the C++ library, *libC.a*, is searched before the HP-RT version of the POSIX library, *libp.a*. Without specifying *-lC*, *CCrt(1)* will direct *ld(1)*, by default, to search *libC.a* after any libraries specified in the runstring. This causes a problem when the POSIX library is needed because there are references to POSIX functions in the *libC.a* routines. By searching *libC.a* before *libp.a*, these symbols are resolved.

Q: I am using the *crt(1)* and *CCrt(1)* scripts to compile and link my HP-RT applications. I know that this will set up searches of the HP-RT versions of the libraries and that HP-UX libraries are also searched. I'm confused about the search order, though. Will you please clarify this for me?

A: The *CCrt(1)* and *crt(1)* scripts simply set up environment variables that are used by *CC(1)* and *cc(1)* to search the correct libraries from both HP-RT and HP-UX for compiling and linking HP-RT applications. The library search order that is defined within *CC(1)* and *cc(1)* has not changed. The linking rules also remain the same. Use the same options for *CCrt(1)* and *crt(1)* as you would for *CC(1)* and *cc(1)*. Let's look at the *CCrt(1)* script for a 2.11 HP-RT Operating system.

```
# cat $HPRTroot/hpux/bin/CCrt
```

```
#!/bin/sh
```

```
# Copyright Hewlett-Packard Co. 1991
```

```
# @(#)HPS1106 B.01.00 HP Embedded Systems Toolkit 09/09/91
```

```
CCINC=/usr/include/CC
```

```
INC=$HPRTroot/usr/include
```



```

ULIB=/usr/lib
LIB=$HPRTroot/lib

HP_CC_REVISION=`/usr/bin/what -s /bin/cc |sed -e s/^[^.]*/[\.]\
([0-9][0-9]\)\.[\]\([0-9][0-9]\)\.*\/\12/' -e 's/0\
([0-9][0-9][0-9]\)\1/' -e '\cc/d'`

FIXED="+A +DA1.1 -DHP_CC_REVISION=$HP_CC_REVISION"

: ${ST_CXXSTDINCL:=$CCINC:$INC}           ; export ST_CXXSTDINCL
: ${ST_CC:=/bin/cc}                       ; export ST_CC
: ${ST_CPP_COMPAT:=$ULIB/cpp}             ; export ST_CPP_COMPAT
: ${ST_CPP_ANSI:=$ULIB/Cpp.ansi}          ; export ST_CPP_ANSI
: ${ST_CCOM:=/lib/ccom}                   ; export ST_CCOM
: ${ST_ASSEM:=/bin/as}                     ; export ST_ASSEM
: ${ST_LD:=/bin/ld}                       ; export ST_LD
: ${ST_PRE:=/lib/pcc_prefix.s}            ; export ST_PRE
: ${ST_LIBC:=$LIB/libc.a}                 ; export ST_LIBC
: ${ST_LIBCL:=$ULIB/pa1.1/libcl.a}        ; export ST_LIBCL
: ${ST_CRT:=$LIB/crt0.o}                   ; export ST_CRT
: ${ST_DLD:=$LIB/libdld.a}                ; export ST_DLD
: ${ST_LIBM:=$LIB/libm.a}                 ; export ST_LIBM
: ${ST_CRTX:=/lib/crtx.o}                 ; export ST_CRTX
: ${ST_LIBH:=$LIB/lib}                    ; export ST_LIBH
: ${ST_LINKCAT:=$LIB/ld.cat}              ; export ST_LINKCAT
: ${ST_CCAT:=$LIB/cc.cat}                 ; export ST_CCAT
: ${ST_ACAT:=$LIB/as_msgs.cat}            ; export ST_ACAT
: ${ST_XDBEND:=$LIB/end.o}                ; export ST_XDBEND
: ${ST_MILLI:=$LIB/milli.a}              ; export ST_MILLI
: ${LD_PXDB:=/usr/bin/pxdb}              ; export LD_PXDB
: ${LPATH:=$ULIB}                         ; export LPATH
: ${CCLIBDIR:=$LIB}                      ; export CCLIBDIR
: ${TMPDIR:=/tmp}                        ; export TMPDIR

if [ $# -eq 0 ]
then exec /usr/bin/CC $FIXED
else exec /usr/bin/CC $FIXED "$@"
fi

```

HP-RT C++ applications using *CCrt(1)* to compile/link will use the following search rules:

1. Any libraries specified in the *CCrt(1)* runstring with

```
-L<directory_path> -l<x>
```

will be searched first.

2. Next, libraries specified in the *CCrt(1)* runstring with

```
-l<x>
```

will search the path specified in the *CCrt(1)* script defined in the *CCLIBDIR* environment variable. (*\$HPRTroot/lib*)

3. The *CCrt(1)* script specifies the environment variable *LPATH* (*/usr/lib*). This sets up the default directory path for library searches if symbols are not resolved by a search of the HP-RT library directory defined by *CCLIBDIR*.
4. You will notice that the *CCrt(1)* script also defines several other environment variables that *CC(1)* will use for defining libraries for the link editor, *ld(1)*, to search. For example:

```

ULIB=/usr/lib
LIB=$HPRTroot/lib
...
: ${ST_LIBC:=$LIB/libc.a}                ;export ST_LIBC
: ${ST_LIBCL:=$ULIB/pa1.1/libcl.a} ;export ST_LIBCL

```

The variables *ST_LIBC* and *ST_LIBCL* define the location that *CC(1)* directs *ld(1)* to use when it is searching *libc.a* and *libcl.a*. So, when the *libcl.a* library is searched, */usr/lib/pa1.1/libcl.a* from HP-UX will be used. When *libc.a* is to be searched, it will use the *\$HPRTroot/lib/libc.a* version.

The search path is very similar when using *crt* to compile/link applications.

1. Any libraries specified in the *crt* runstring with

```
-L<directory_path> -l<x>
```

will be searched first.

2. Next, libraries specified in the *crt* runstring with

```
-l<x>
```

will search the *\$HPRTroot/lib* directory. This is because the script sets the *LPATH* environment variable, making this the default directory for

-l<x> library searches.

3. The *cc(1)* compiler will similarly use environment variables to determine where to search for other default libraries. It will use a predetermined order defined by *cc(1)*.

You can use the *-v* option in the compile runstring to display the actual order that *CC(1)* and *cc(1)* instruct *ld(1)* to use.

Q: I have a PCMCIA flashdisk that I would like to boot a Ramdisk-based system from. I want to use this for running diagnostics and I do not want networking set up. I would like it to be strictly memory-based. Is this possible?

A: Yes, this is possible.

1. Make a customized copy of the *\$HPRTroot/etc/conf/ramdisk.def* file to include only the programs you want in your Ramdisk-based system. Use the guidelines in the *ramdisk.def* file itself to determine the minimum file set and add programs as desired.
2. Modify the *\$HPRTroot/etc/ramdisk/rc* file so that it does not try to execute *Mount_Target* or any other programs that no longer exist. This will eliminate network setup.
3. Build your new ramdisk kernel.

```
# make -f $HPRTroot/etc/conf/hp-rt.make CPU=s743 PCMCIA= RAM= \
    RAMDEF=$HPRTroot/etc/conf/<custom_ramdisk_filename> \
    KERNEL_FILE=<your_kernel_file_name>
```

4. Configure and boot a disk-based PCMCIA kernel. See the *HP Z5117A PCMCIA Adapter Installation and User's Guide*, Chapter 4, for information on how this is done.
5. FTP your Ramdisk kernel to the PCMCIA disk-based file system. Remember to change the file permissions after you have transferred the file to the disk.
6. Boot your new Ramdisk-based kernel:
 - a. Halt your 743rt system:

```
# /etc/reboot -h
```

- b. Reset the 743rt.
- c. Select the PCMCIA flashdisk as your boot device.
- d. From the ISL prompt:

```
ISL> rtboot -a -rramdisk /your_kernel_file_name
```

You should now be running your Ramdisk-based system.

Q: I want to install HP-RT Revision 2.2 into directory */HP-RT* on my HP-UX 10.10 host. It will only let me install to */opt/HP-RT*. Am I no longer able to install to */HP-RT*?

A: It is no longer possible to install directly into the HP-UX root directory when using */usr/sbin/swinstall*. What you need to do is create a symbolic link pointing to */HP-RT* from */opt/HP-RT*:

```
# ln -s /opt/HP-RT /HP-RT
```

Then you can run *swinstall(1M)* and install the HP-RT bundle. It will install to */opt/HP-RT* but you can access it using */HP-RT* the same as you did with previous revisions.

Q: I want to install HP-RT software into multiple trees on the same system. I don't see how to change the target location using *swinstall*. Can this still be done?

A: Yes, you can still install HP-RT software into multiple locations using */usr/sbin/swinstall*. Assume, for example, that you have already installed software onto the */opt/HP-RT* tree and you now want to install it onto */HP-RT_2.2*.

1. Log on as root.
2. Run */usr/sbin/swinstall*.
3. From the 'Options' menu, select 'Change Options...'
4. Mark the 'Allow creation of

multiple versions' selection.

5. Open the software bundle you wish to install:

B5487AA -> A.02.20 HP-RT Developer's Kit

6. Select the product to install:

HPRT-OS-PRD -> A.02.20 HP-RT Core Product

7. From the 'Actions' menu, 'Mark for install'

8. Select the product again:

HPRT-OS-PRD -> A.02.20 HP-RT Core Product

9. From the 'Actions' menu, 'Change Product Location'

10. Mark the 'Other Location' selection

11. Enter the parent directory path where you wish to install the HP-RT root tree:

Path: /HP-RT_2.2

12. Select 'OK'

The software will be installed into the directory /HP-RT_2.2 for our example, but you will see by investigating the /var/adm/sw/swagent.log file that it has not been configured. You will need to run swconfig separately for each product you are installing into /HP-RT_2.2.

Do this as follows:

```
# swconfig -x allow_multiple_versions=true <product_name>, \
    l=<parent_directory_path>, \
    r=<revision>
```

For our example, do:

```
# swconfig -x allow_multiple_versions=true HPRT-OS-PRD, \
    l=/HP-RT_2.2,r=A.02.20
```

You can see by checking the /var/adm/sw/swagent.log that your software has now been installed and configured into the /HP-RT_2.2/HP-RT directory. Check /var/adm/sw/swconfig.log for information on what swconfig has done.

Using this method allows you to track software installation by looking at the /var/adm/sw/swagent.log file. swinstall will make entries in this file to show what has been installed on your system, where, and when.

Q: I want to create a recovery system that will work with a graphics console. When I try to boot the one I created following the directions

in the *HP-RT System Administration Tasks* manual, Chapter 4, "Making a Recovery System," it starts to boot on the graphics console but then the output disappears. I assume that it reverts to the RS232-port A. What do I need to do?

A: You need to modify \$HPRTroot/etc/conf/recovery/make_recover to add GRAPH3= to the hp-rt.make make options in the file as follows:

```
MAKE_OPTIONS="RAM= NONFS= NONFSS= NOPTY= NOLAN0=
NOTCPIP= GRAPH3="
```

This will make the recovery system include the X11 SERVERrt device drivers in the kernel file so that the recovery kernel can communicate with the graphics console.

Also, if you are going to use a graphics console to build an install kernel and a disk-based system, you will need to make similar modifications to the \$HPRTroot/hpux/adm/BuildKernels file.

In the "Building install kernel..." section, add the GRAPH3= option to the make command:

```
/bin/make -f $CONF/$FILE.make RAM= KERNEL_FILE=
$INST_KRNL \
CPU=$CPUTYPE RAMDEF=$INST_DEF GRAPH3= $KERNpcmcia
```

In the "Build the disk-based kernel..." section, add the GRAPH3= option to the make command:

```
/bin/make -f $CONF/$FILE.make CPU=$CPUTYPE RDB=
$KERNpcmcia \
$KERNbpn $KERNstrms $KERNx11srv $KERNsnmp GRAPH3= \
KERNEL_FILE=$DISK_KRNL
```

Then you will be able to install a disk-based system using your graphics console. ■

HP-RT Operating System questions are answered by Anita Harris, a support engineer in the HP-RT Worldwide Technology Expert Center. She has worked with Real Time systems for 12 years as a customer, a Real Time Response Center Engineer, and HP-RT WTEC Engineer. She can be reached at harrisa@mayfield.hp.com.



New Products

Multiplatform Programmer's Editor

Vital Inc. has released CRiSP Version 5, a cross-platform programming text editor for X Windows, Windows, and character-based environments. The user interface provides easier navigation for both programmers and casual users. Complete documentation is now available online in hypertext format on all UNIX and Windows platforms.

CRiSP v5 integration with multiple source code control systems provides integrated check in/check out facility from within, resulting in a highly integrated development environment. Users can now create and paste language-specific templates visually from within the editor. It is now easier to add colorizer definitions for specialized languages, as well. Dynamic colorization and printing is supported for over 50 languages, including Java and HTML.

Users can now perform searching/replacing and other advanced arithmetic operations within columnar ranges. CRiSP v5 is completely encapsulated with HP Softbench and other development environments. Keyboard emulation is provided for Wordstar, BRIEF, vi, and CUA.

CRiSP v5 is available on HP-UX and other platforms. It is priced starting at \$149.99 on Linux, \$249.99 on Windows, and \$349.99 on UNIX.

Contact Vital Inc., phone: (214) 491-6907, fax: (214) 491-6909, e-mail: info@vital.com, URL: <http://www.vital.com>.

X Window Reference

Digital Press has announced *X Window Systems: Core Library & Standards*, by Robert W. Scheifler and James Gettys. The first volume of the updated reference includes new information on X

Version 11, Release 6.1.

Authored by the experts who originally designed and implemented the X Window system and with new, significant input from X Window Consortium members who have helped evolve X Version 11, Release 6, this fourth edition is a major revision of this definitive reference. Updated to conform with the latest released software (X Version 11, Release 6.1), this edition also has been reorganized into three separate books to make each more usable for its targeted audience. Each volume provides its own comprehensive subject-oriented index that further enhances its overall value.

X Window Systems: Core Library & Standards, a 400-page paperback (ISBN: 1-55558-154-4), is \$49.95.

Contact Digital Press, fax: (617) 928-2620.

COBOL Application Migration

Unidata, Inc. has announced Version 2.1 of COBOL Direct Connect, designed to extend the life and utility of existing legacy COBOL applications.

COBOL Direct Connect brings Unidata's "nested" relational database management system to existing COBOL applications. It can migrate COBOL applications and associated indexed file structures to open systems without significant data file restructuring, code reengineering, or performance penalty.

The product currently supports COBOL code running RM/COBOL, Micro Focus COBOL, and Acucobol. It is available on HP-UX and other UNIX systems.

Unidata estimates that companies choosing an incremental migration path revolving around COBOL Direct Connect can achieve performance improvements of 4 to 10 times over

Alex meets DLT...

and they back up at full speed

Now IEM offers Alexandria Network Backup Librarian with high-speed DLT libraries, so you can relax while your network backs up FAST! With DLT4000 drives, the single drive transfer rate is up to 3 MB/sec. with 2:1 compression; with 9 drives, the library throughput is up to 97.2 GB/hour! DLT4000 cartridge capacity is 40 GBytes with 2:1 compression.


Alex backs up to DLT libraries with 2 to 9 DLT4000 drives and 28 to 264 cartridges.

These libraries are so reliable that they include a full one-year on-site warranty with next business day response.

Alex also backs up to IEM's 4mm and 8mm carousels and autofeeders, as well as 3480 tape libraries and HP's optical libraries.

Alex supervises every facet of your heterogeneous UNIX network backup:

- Organizes and automates your system backups
- Manages all of your backup media
- Performs cold database backups of Oracle, Sybase and Informix databases
- Performs hot backups of Oracle databases with an additional module



IEM: Providing Solutions for a Lights-out Environment



In the U.S. and Canada:
IEM, Inc., P.O. Box 1889
Fort Collins, CO 80522 USA
Phone: (970) 221-3005
(800) 321-4671
Fax: (970) 221-1909

In the United Kingdom:
IEM, Inc., Unit 6, Salisbury House,
Wheatfield Way, Hinckley Fields,
Hinckley, Leicester LE10 1YG
Phone: +44 (0)1455 239000
Fax: +44 (0)1455 239668

All others:
IEM International Sales
1629 Blue Spruce Drive
Fort Collins, CO 80524 USA
Phone: +[1] 970-221-3005
Fax: +[1] 970-221-1909

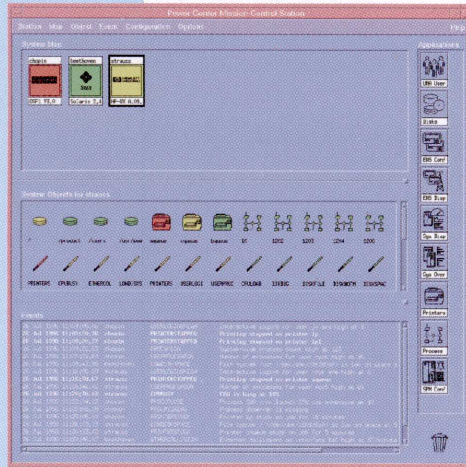
Enterprise System Management

Power Center Software has announced Power Center 2.0 through its "I Lov It" Program, through which qualifying sites will receive a permanent license for the Power Center System Problem Monitor agent for each Sun Solaris, HP-UX, and/or Digital UNIX system they are currently operating. They also receive the Mission Control Station (MCS) GUI console for their enterprise. These two products provide a complete solution for effective enterprise system management.

"I Lov It" customers pay only for value added past the base solution. They may license additional products from the Power Center suite or purchase software support and/or update services. Product and service pricing is permanently guaranteed to "I Lov It" customers. Each application in the Power Center suite carries a flat price for both servers and clients, and the package pricing of the agent allows users to expand their enterprise without any penalty.

Power Center 2.0 includes additions to MCS console, which now includes a system map, system object display, event console, and application toolbox. It displays the status of systems and their components through color icons and event messages from across the entire enterprise.

Contact Power Center Software LLC, phone: (303) 220-9789.



Power Center 2.0

Oracle and to link multiple tables into a single file for subsequent reporting or processing. A new Suprtool module, STExport, lets users prepare data for export from HP-UX to virtually any other application or operating platform, without requiring users to write custom programs.

With Suprtool/UX, Oracle users benefit

from the familiar SQL syntax for Oracle selections, while users migrating from MPE can use the same commands in HP-UX as in Suprtool/MPE. Users can choose between SQL or Suprtool sort commands, using the one that gives the best performance.

With STExport, users can now prepare data in formats that can be imported into spreadsheets, MS Access, and Oracle on such platforms as PCs and UNIX.

Suprtool/UX 3.8 is being shipped to all users who have an active support contract. Robelle charges the same price for all models of HP-UX processor.

Contact Robelle, phone: (800) 561-8311 or (604) 582-1700, e-mail: info@robelle.com.

Online Backup for Informix Databases

Legato Systems, Inc. has announced its NetWorker Database Module for Informix, which provides storage management services for INFORMIX-OnLine Dynamic Server's enhanced backup and restore utility. This allows database administrators to back up Informix databases while they remain online

and in use. Legato's Database Module allows frequent, regularly scheduled "lights-out" database backups to ensure high system availability.

The Database Module features fully unattended, "lights-out" operation and includes the ability to schedule partial or full backups on a regular or rotating schedule. NetWorker also offers support for a wide range of optical and tape devices, including autochangers, which minimizes operator intervention.

Legato's Database Module for Informix was scheduled to begin shipping in September 1996 and will be available on HP-UX and other UNIX systems.

Contact Legato, phone: (415) 812-6000, fax: (415) 812-6032, faxworker: (415) 812-6156, URL: <http://www.legato.com>.

Book on UNIX Clusters

Prentice Hall PTR and Hewlett-Packard Professional Books have announced *Clusters for High Availability* by Peter S. Weygant, the complete guide to using clusters of UNIX computers to gain high levels of availability and HP's cluster solutions at work.

Clusters for High Availability explains the basic concepts, architectures, and terminology of high availability using clusters and gives a broad outline of how HP's cluster solutions work. The book describes

what could be achieved by applications rewritten with embedded SQL. Once stored in UniData RDBMS, legacy COBOL applications can benefit from a rich application development environment and database interoperability products.

Contact Unidata, phone: (303) 294-0800, fax: (303) 293-8880, e-mail: unidata@unidata.com.

Data Access

Robelle Consulting Ltd. has announced Version 3.8 of the Suprtool data access utility for HP-UX. Suprtool/UX can now access Oracle views, in addition to reading and writing data files with fixed-length records. Suprtool/UX now allows users to select and sort data from

ISAMATION

THE FASTEST ISAM REBUILD AND DATA RECOVERY TOOL AVAILABLE!

HP-UX COBOL

TISAM

Micro Focus COBOL

Informix S.E.

C-ISAM

For more information on the ultimate ISAM maintenance and data recovery tools, contact:



Computer Solutions, Inc.

120 E. Marks St. • Suite 225 • Orlando, FL 32803
407-649-0123 or 512-343-6634 • FAX 407-649-1407

All trademarks belong to their respective holders.

CIRCLE 111 ON READER SERVICE CARD

how high-availability clusters can be configured to eliminate single points of failure in power systems, disks, processors, network components, and software.

The book features coverage of HP high-availability products and presents suggestions for developing high-availability solutions in publishing, mail order catalogue sales, brokerage services, insurance, and other industries.

Clusters for High Availability, by Peter Weygant (ISBN 0-13-494758-4), is priced at \$27.

Contact Prentice Hall PTR, phone: (201) 236-7139, fax: (201) 236-7131.

Free Java Library Use

ObjectSpace, Inc. has announced that the Java Generic Library (JGL) is now available for free commercial use. The Java Generic Library is designed specifically for performance and use in a distributed environment and is completely compatible with Sun Micro-

systems' Java Developer's Kit (JDK).

ObjectSpace has augmented the Java Developer's Kit by leveraging technology from ObjectSpace C++ toolkits and other vendors' libraries to form the Java Generic Library. The JGL algorithms can be used independently of the JGL containers and can be applied to native Java arrays of objects, native Java arrays of primitives, all Java Generic Library containers, and *java.util*. vectors from the existing Java Developer's Kit. This enables the JGL to remain small and efficient.

The standard Java Generic Library release includes full source code, online HTML documentation, comprehensive examples, a 100-page tutorial, and a suite of performance benchmarks.

The Java Generic Library is free and available from the ObjectSpace Web site (<http://www.objectspace.com>).

Contact ObjectSpace, phone: (214) 934-2496, fax: (214) 663-9100, e-mail: jgl@objectspace.com.

Healthcare and Process Manufacturing

Avnet Computer and Ross Systems have jointly announced the CONTINUUM/TS integrated solution targeted at healthcare and process manufacturing organizations. The joint program is an outcome of the Avnet Computer/Ross Systems Alliance, formed earlier this year.

Designed to meet financial, technical, and process manufacturing needs, the CONTINUUM/TS solution features high-performance, open systems hardware from either DEC, HP, or IBM; Renaissance CS client-server software from Ross Systems; and Avnet Computer's technical and financial consulting services to both manage and execute technology platform migrations. Through CONTINUUM/TS, customers can acquire computer systems, workstations, peripherals, or entire networks and exchange them for new technology when it becomes available. The cost is


RAID Storage

ANDATACO has introduced ESP RAID Lite, which combines the fault-tolerant advantages of RAID with Enterprise Storage Packaging (ESP) technology. ESP RAID Lite embodies single active or dual active/passive controllers, each of which employs an i960 Intel RISC microprocessor. This processor provides visual, audible, e-mail, and Web-based alarms preceding failure of a system's disk drives, fans, or power supplies. Once warned, system administrators can repair a faulty component before any damage occurs and significantly reduce their system's downtime.

ESP RAID Lite's microprocessor supplies the needed MIPS to provide I/Os between the fastest RISC-based CPUs and disk drives available. ESP RAID Lite can thus facilitate imaging, video-on-demand, bulk storage, and multiple online transactions. ESP can operate UltraSCSI (40 MB/s) data transfer rates and support 10,000 RPM drives.

Contact ANDATACO, phone: (800) 334-9191 or (619) 453-9191, fax: (619) 453-9294, e-mail: inquire@andataco.com.

ANDATACO ESP Storage



a fixed monthly fee that is generally less than current operating expenses.

Contact Avnet Computer, phone: (602) 414-3304, fax: (602) 414-5507.

Technical Desktop Printing

AutoGraph International (AGI) has announced a suite of imaging and printing tools for workstations.

With EasyCopy/X 5.0 a user captures and prints screen images and raster files to raster, PCL, RTL, and PostScript printers. A new, optional module, EasyCopy/Graphics, enables users to load, display, and print the most commonly used vector files. It prints to the raster and/or PostScript printers currently in use.

The new EasyCopy/X also features a direct link to With EasyCopy/Scan, a new scanning solution. The module functions as a stand-alone program or it may be installed as an integrated part of either With EasyCopy/X or Flexprint. Another option, EasyConvert, allows for quick file conversion of raster formats. EasyCopy/NT supports output of screen captures and selected file formats to raster and PostScript printers by means of the Windows drivers. Flexprint 2.0, an image printing system for PostScript

level 2 printers, intelligently interprets the printer's options and features while ensuring the automatic support for future PostScript printers.

The AGI suite of printing and imaging tools is available on CD-ROM. The products will run in demo mode until a license is purchased. EasyCopy/X pricing begins at \$395, while Flexprint starts at \$695 for single-user licenses. The companion programs are individually priced. EasyCopy/Scan is priced at \$695 for a single user; EasyCopy/NT licensing is \$149 for a single user; EasyConvert and FlexImage pricing starts at \$195.

Contact AutoGraph International, phone: (408) 436-7227, fax: (408) 436-7255, URL: <http://www.augrin.com>.

Enterprise Client-Server Development

Andersen Consulting has announced Version 3.x of FOUNDATION for Cooperative Processing (FCP), which features new OpenRamp technology that easily integrates popular workgroup applications and development tools into the FCP enterprise architecture. FCP, which consists of a high-performance, open architecture and toolset, is a widely

used enterprise client-server application development environment.

FCP OpenRamp enables developers on Windows NT to use OLE-based workgroup development tools, such as Visual Basic or PowerBuilder, to build interfaces and client-side logic for FCP-based enterprise systems to leverage existing client-side skills and to migrate existing applications built with these tools into the FCP enterprise architecture.

A fully configured FCP 3.x developer toolkit software license, which includes OpenRamp, is priced at \$10,800. Applications built with FCP can be easily generated to run on HP-UX and other server platforms.

Contact Andersen Consulting, phone: (847) 714-4000.

Network Connectivity

Specialix Inc. has announced the new JETSTREAM products for WAN/LAN, TCP/IP, and Remote Access Servers. The JETSTREAM line expands Specialix's terminal server product range, adding cost-effective (and improved) network connectivity.

The JETSTREAM 6000 is an 8-port terminal server that is upgradable to 16, 24, and 32 ports, with the use of clip-on modules. JETSTREAM 6000 supports

rologin/telnet, multisessions, downloadable terminal definitions, and printer or modem grouping.

JETSTREAM 7000 and 7500 are expandable from 8 to 16 ports and provide a single point of entry for remote access and WAN applications such as LAN bridging, home and remote office, and Internet and database access. The products support SLIP/PPP, 115.2-Kbps line speed, and hunt group capability, as well as SNMP (MIB II), BOOTP, net rebooting, and downloadable terminal definitions.

JETSTREAM 6000 is priced at \$1,695. JETSTREAM 7000 is priced at \$1,795. JETSTREAM 7500 is priced at \$2,395.

Contact Specialix, phone: (800) 423-5364 or (408) 378-7919, e-mail: info@specialix.com.

Expert Systems Tool

Production Systems Technologies, Inc. has announced CLIPS/R2, the first implementation of CLIPS to use the proprietary Rete II rule engine. CLIPS/R2 is said to be more than 50 times faster than the previous version on complex problems.

CLIPS/R2 is upward-compatible with the most recent previous version of CLIPS, CLIPS 6.05. CLIPS/R2 supports forward-chaining rules, backward-chaining rules, objects, and conventional procedures. The new backward-chaining rule engine in CLIPS/R2 supports certainty factors, ANDs and ORs in rules, and HOW and WHY explanations. CLIPS/R2 rules bases can be embedded in C or C++ programs. The new CLIPS Interface Definition Compiler (CIDC) enables the CLIPS/R2 rules to operate transparently on existing C structs and C++ classes.

CLIPS was developed by NASA

Johnson Space Center in 1985 and has become one of the most widely used expert systems tools in the world.

CLIPS/R2 is available for Windows and UNIX systems.

Contact Production Systems Technologies, phone: (412) 683-4000, fax: (412) 683-6347.

New from PLATINUM technology

HP OpenView IT/Operations Integration

PLATINUM technology has announced that the PLATINUM Apriori problem resolution products are fully integrated with the Hewlett-Packard OpenView IT/Operations management framework. Apriori recently completed the HP IT/Operations certification program and has achieved the Premier level (highest technical level) of integration with IT/Operations.

IT/Operations managers can now use

the Internet and corporate intranets and the IT/Operations console to quickly solve problems via the Apriori Hands-Free Support tool. IT/Operations managers can also launch and manage Apriori, send and receive messages/alerts to and from Apriori, and pass information between Apriori and other systems management tools.

Oracle Data Unloading

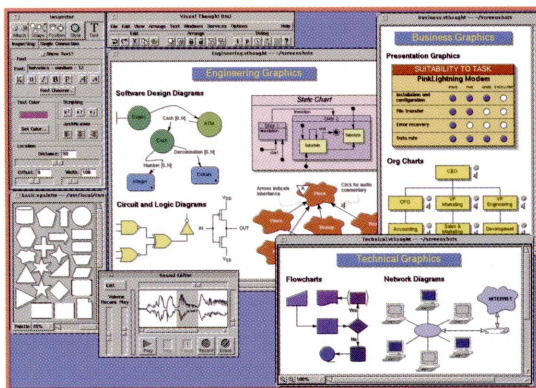
PLATINUM technology has also announced PLATINUM Fast Unload for Oracle, a software utility that enables database administrators (DBAs) to extract data from an Oracle database table up to 23 times faster than Oracle Export.

Fast Unload achieves its high speed by directly accessing Oracle database files, completely bypassing the SQL engine and Oracle kernel. No other data unloading product accesses Oracle database files in this manner.

Fast Unload offers selective data

visual thought

UNIX Diagramming & Flowcharting



Use Visual Thought for:

- Software design diagrams (Booch, Rumbaugh, Objectory, Fusion, custom notations)
- Flowcharts for ISO 9000, TQM, BPR
- Clickable Web diagrams with GIF/JPEG and server/client-side imagemap export
- Framemaker documentation graphics w/ MIF, EPSI export
- Network diagrams with network component clip art
- Dataflow diagrams, org charts
- Presentation & business graphics
- Graphical pre/postprocessor for other tools or simulators

FREE CD-ROM, FTP & 30-day trial

Additional benefits

- Intelligent, rubberbanding lines
- Macintosh-like ease-of-use... on UNIX!
- Editable, WYSIWYG drag-and-drop palettes
- Hyperlinked, hierarchical documents
- Arbitrary object rotation
- Complete text handling: subscripts, arbitrary fonts, sizes, colors, styles and justification
- Dozens of export formats: GIF89, JPEG, MIF, TIFF, EPSI, XWD, SunRaster, others
- Available on SunOS, Solaris, HP-UX, and soon, Windows 95/NT

VISIT (use code AE14)
<http://www.confluent.com/>
 800-780-2838 ext. 154
 info@confluent.com

©1996 Confluent, Inc. All rights reserved. Visual Thought, the Visual Thought logo, Confluent, and the Confluent logo are trademarks or registered trademarks of Confluent, Inc. All other names are the properties of their respective holders.



Confluent, Inc.
 132 Encline Court
 San Francisco,
 California 94127
 415-586-8700 Voice
 415-586-8838 Fax

See why Visual Thought has thousands of users worldwide! Clients include:

ABB • AT&T • Alcatel • BNR • Bellcore • Boeing • CS First Boston • Credit Suisse
 E-Systems • Eds • Ericsson • Ford • HP • IONA • Lucent • McDonnell Douglas
 Motorola • Nokia • Nemura • Nortel • Persistence • Prediction
 Pure Software • SAIC • Smith Barney • Sprint • Toshiba
 U.S. Air Force • U.S. Navy • U.S. West • UBS
 Wells Fargo • Nikko

CIRCLE 37 ON READER SERVICE CARD

unloading, allowing users to extract specific rows and columns from a tablespace, further reducing the time it takes to unload select data; flexible output formatting options, allowing output records to be automatically formatted for compatibility with PLATINUM Fast Load for other database load utilities; and multiple output files, allowing for unloading of a table that is larger than 2 GB (the maximum file system size on most UNIX systems today), increasing the reliability of an unload process.

Database Administration

PLATINUM technology has also announced Desktop DBA for UNIX, which replaces the need to memorize SQL syntax for hundreds of RDBMS commands with point-and-click commands.

Routine and complex manage-

ment tasks for multiple database types that are supported include automating database comparisons, saving administrators' time when synchronizing databases; altering database objects—including tables, indexes, views, and procedures—in real time simply by clicking on dialogue box options; the ability to drag-and-drop objects and data—including all dependencies and associated security—between database servers for efficient migrations; and streamlining administrative functions, such as those for local and remote server configuration, user access permissions, and catalogue reporting.

Prices start at \$2,490.

Contact PLATINUM, phone: (800) 442-6861 or (708) 620-5000, e-mail: info@platinum.com.

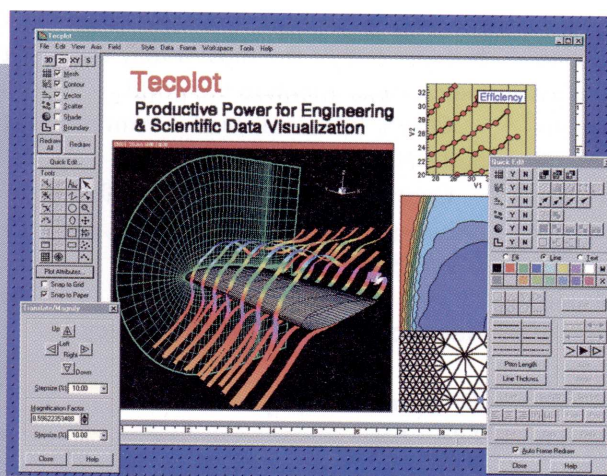
Data Visualization

Amtec Engineering, Inc. has announced Tecplot Version 7.0, a set of tools available for visualizing and plotting large amounts of data. This release features a new GUI, animation, and page layout. Tecplot helps users work productively with the datasets generated by numerical simulation (such as computational fluid dynamics), statistical analysis, data acquisition, and other sources.

Users can annotate plots and adjust the amount of information in a plot to match the audience's familiarity with the data.

Tecplot Version 7.0 runs on most UNIX workstations (under Motif) and PCs running Microsoft Windows (3.x, 95, and NT). Single-user pricing ranges from \$995 to \$3,195, depending on platform and license type. Amtec offers a 60-day, money-back guarantee.

Contact Amtec Engineering, phone: (206) 827-3304, URL: <http://www.amtec.com>.



Amtec Engineering Tecplot

New from Artecon

8-mm Tape Drive Subsystems

Artecon has announced the Lynx 40-GB 8-mm tape drive. The follow-on product to the Lynx 14-GB 8-mm tape drive, this drive is fully read backward compatible.

The transfer speed and cache have been redesigned for better performance and efficiency. The transfer speed of the 40-GB tape drive at 6 MB/sec (compressed) is double that of the 14-GB drive. Cache has also increased from 500 KB to 4 MB.

The Lynx 8-mm tape subsystem includes modular stackability, external SCSI and AC jumpers for reducing cable tangles and SCSI bus length, internal PCBs to increase overall subsystem reliability by eliminating cable connections, and externally removable fans.

Price for the Lynx 40-GB 8-mm tape drive is \$7,295.

DLT Backup Subsystems

Artecon has announced a full line of 19-inch EIA rackmountable DLT. The LynxRack DLT complements the LynxRAID+ controller series and Lynx mass storage units. This subsystem offers many high availability features found in RAID controller subsystems.

The LynxRack DLT supports two front removable, load-sharing power supplies that provide N+1 redundant power dedicated to each drive. Dual removable fan modules provide cooling to the power supplies and DLT tape drives. Dual AC outlets provide additional redundancy for the subsystem.

Structural Research & Analysis COSMOS/Works

The DLT-4000 holds 20 GB native and up to 40 GB compressed, while the DLT-2000XT holds 15 GB native and up to 30 GB compressed. The LynxRack DLT does not impede the fast 3.0 MB/sec compressed transfer speed of either drive.

Prices for the LynxRack DLT vary depending on the type and number of DLT drives and start at \$6,995.

Contact Artecon, phone: (619) 931-5500.


Secure Telnet Access

WRQ has announced Reflection Secure 5.2, a new version of the company's software that implements the Kerberos V5 protocol for secure end-to-user authentication. The new version supports Windows 95 and Windows NT. When running on a PC with other Reflection host connectivity products, the software provides secure Telnet access to UNIX, Digital, and HP host computers in a networked environment. As a client of Kerberos V5 or DCE security servers, Reflection Secure eliminates the need to send passwords over a network and protects against unauthorized use.

The software features secure single sign-on, user password changes, and ease of use. Automatic credentials acquisition from the Reflection Telnet session minimizes the need for user training. Advanced options for credentials protection on Windows clients and support for multiple credentials and multiple principals are provided.

Reflection Secure 5.2 is priced at \$99.

Contact WRQ, phone: (800) 872-2829 or (206) 217-7100, fax: (206) 217-0293, e-mail: info@wrq.com.



Report Viewing Tool

New Dimension Software has announced Version 2.10 of its CONTROL-D/PC for Windows report viewing tool. The tool enables users to download reports from New Dimension Software's CONTROL-D mainframe report distribution software or CONTROL-V mainframe report archiving software. Using CONTROL-D/PC installed on a Windows-based PC, these mainframe reports can be easily accessed for storage, viewing, modification, or printing. Reports can be archived to hard disk drives and removable media such as optical disks and CD-ROM. Other key features include a new GUI for the CONTROL-D/Delivery report distribution solution. Additionally, Version 2.10 now can import and view PostScript and Adobe Acrobat files.

CONTROL-D/PC can be purchased as a companion product for CONTROL-D and CONTROL-V. Version 2.10 is priced at \$250 per seat.

Contact New Dimension Software, phone: (800) 347-4694, ext. 522 (North America), or 972-3-645-1111 (Israel), URL: <http://www.ndsoft.com>.

Backup for Sybase, Oracle

PDC has announced that its BudTool backup management tool has added support for Sybase 11 and Oracle 7.x. BudTool Support for Sybase works in conjunction with BudTool and uses the Sybase API and backup server to assure

Design Analysis

Structural Research & Analysis Corporation has announced the addition of direct model transfer to Phase I of its COSMOS/Works analysis interface to SolidWorks. Users will no longer need to save their SolidWorks models as IGES files for transfer to COSMOS/Works. Benefits include faster analysis, with no delay during model transfer; greater accuracy, as no information will be lost during translation; easier, more successful meshing of models; and less internal bookkeeping due to duplication of information.

COSMOS/Works gives SolidWorks users the design analysis capabilities of COSMOS/M and COSMOS/FFE (Fast Finite Element). A powerful flowchart guide walks users through the analysis process while answering "what if?" questions. COSMOS/FFE solution technology solves complex problems up to 100 times faster than conventional analysis codes while reducing disk storage space up to 95 percent, enabling users to solve mainframe-sized problems on low-cost engineering workstations, PCs, and laptops, the company notes.

Contact Structural Research & Analysis, phone: (310) 207-2800, ext. 770, e-mail: louise@srac.com.

database integrity and performance. BudTool Support for Oracle provides full autodiscovery at the datafile level in a tablespace. The new database modules are supported on the Sun Solaris and HP-UX platforms.

BudTool Support for Sybase permits hot backup of databases and transaction logs. Sybase database administrators can schedule automated backups and simplify retrievals using BudTool's sophisticated media and automated tape library capabilities. Data can be restored from backup media with or without BudTool installed on the network.

Oracle database backups and restores are scheduled and performed using the

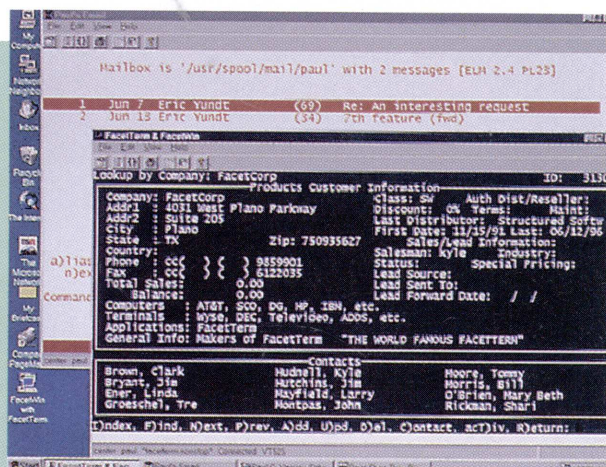
Windows-to-UNIX Connectivity

FacetCorp has announced FacetWin, which allows Windows 95 and Windows NT users to transparently access and use any UNIX-based network resources (such as files, disks, applications, and printers). Features include transparent file and print services, terminal emulation, the ability to automatically back up networked PCs to a UNIX tape drive, networked modem access for PC users, an e-mail POP3 server, and remote computing support.

FacetWin's transparent file and print services are based upon Microsoft's Server Message Block (SMB) technology for transferring data. Most existing connectivity products use the Network File System (NFS) method of loading NFS software on the PC to enable it as a UNIX NFS client. FacetWin's UNIX SMB server requires no additional software to be loaded on the PC. SMB technology also provides high throughput for file and print services and is easy to install and administer.

FacetWin is priced at \$195 per user, with discounts for multiuser licenses. FacetCorp is offering free 30-day evaluation copies of the product through worldwide distributors.

Contact FacetCorp, phone: (214) 985-9901, fax: (214) 612-2035, URL: <http://www.facetcorp.com>.



working and scalability.

HP-UX 10.20 includes several features that support 64-bit processor systems. One is large-file—up to 128 GB—support, including JFS, LVM, and high availability. This, in addition to the large-file system support introduced with HP-UX 10.10, gives HP-UX 10.20 64-bit functionality in a 32-bit operating system. HP-UX 10.20's maximum addressable RAM size is 3.75 GB.

HP-UX 10.20 also includes advanced compiler functionality, fast procedure calls, virtual memory mapping, and large-page capabilities. Each of these features allows HP-UX 10.20 to facilitate high-end OLTP, decision support, and technical applications—all while multiple applications are using the system environment.

HP-UX 10.20 is scalable across the entire HP 9000 product line. As with preceding versions of HP-UX, HP-UX 10.20 is tuned to run on SMP systems. With SMP, the operating environment provides excellent scaling of application performance across multiple processors using a single version of the operating system. For enterprise parallel computing, HP-UX 10.20 allows running of parallel applications software for potential gains in performance.

As with previous releases of HP-UX, HP-UX 10.20 provides forward binary compatibility when upgrading to new HP-UX versions and new servers and workstations. Any existing 9.X or 10.X application can be run unmodified in the 10.20 environment. HP-UX 10.20 is an X/Open UNIX 95-branded product, signifying that it conforms to X/Open's Single UNIX Specification.

A two-user license for the HP-UX 10.20 environment is bundled with HP

same screens and interfaces used by BudTool for file system operations. The new database support module permits hot or cold backup of tablespaces, archive logs, and configuration files. Oracle 7.x backup tapes created with BudTool can be restored using standard UNIX and Oracle commands and procedures.

Contact PDC, phone: (800) 654-4732 or (610) 265-3300.

New from HP

HP-UX 10.20

HP has announced its newest, most advanced enterprise-class HP-UX operating system—HP-UX 10.20. This latest version extends HP-UX applicability across the entire HP 9000 product line of enterprise servers, workstations, and

Enterprise Parallel Servers (EPS). HP-UX Version 10.20 is designed to optimize the performance of HP's 64-bit PA-8000 RISC processor.

HP-UX 10.20 offers the following new technical advancements: full support for 64-bit PA-8000 systems; large file size of up to 128 GB for support for Hierarchical and Journaled File Systems (HFS and JFS) and Logical Volume Manager (LVM); support for up to 2 billion user and group IDs; full X11R6 graphics libraries support (client and server); incorporation of a DMAPI (Data Management API) for hierarchical storage management in online JFS (OLJFS); enhancements to Process Resource Manager that allow allocation of RAM by user, group, or process; and symmetric multiprocessing (SMP) support for net-

9000 workstations and servers, with additional license levels available.

SNA Connectivity

HP has introduced new Systems Network Architecture (SNA) interconnectivity solutions. These networking products enable users to more rapidly and cost-effectively link HP 9000 commercial computing environments with large production systems from IBM.

The latest release of HP's HP-UX SNAplus2 features five comprehensive solutions that enhance system scalability and application availability in data-warehousing and mainframe-migration environments. SNAplus2 Link provides enhanced communication capabilities between an HP 9000 server or workstation and an IBM mainframe or peer system, such as an AS/400. HP-UX SNAplus2 3270/3179G provides interactive communications between an HP 9000 server and IBM-compatible mainframe or peer system using SNA 3270, 3278, and 3179G/3192G emulation. HP-UX SNAplus2 APPN End Node allows an HP 9000 system to take full advantage of Advanced Peer-to-Peer Networking (APPN) facilities while still retaining the capability to operate in traditional mainframe host-managed networks. HP-UX SNAplus2 API allows a wide variety of VPIs to enable application communications. HP-UX SNAplus2 Remote Job Entry provides batch data transfer between an HP 9000 system and an IBM mainframe host in an SNA environment.

The HP-UX SNAplus2 portfolio starts at \$3,550 for Tier 1 systems. Stand-alone SNAplus2 starting prices are as follows: HP-UX SNAplus2 Link, \$500; HP-UX SNAplus2 APPN EN, \$600; HP-UX SNAplus2 3270/3179G, \$850; HP-UX

SNAplus2 API, \$300; and HP-UX SNAplus2 RJE, \$1,300.

Intranet Security Software

HP recently unveiled the Praesidium/Security Service for user authentication, and the comprehensive Virtual Vault security solution. With the Praesidium/Security Service, non-DCE users can have DCE-like authentication without a DCE/9000 cell. HP VirtualVault is the first comprehensive security solution that enables electronic transactions end-to-end without compromising the integrity of internal computing environments. HP VirtualVault software, part of HP's Praesidium enterprise security framework, grants internal and external clients Web access to critical applications while safeguarding proprietary information assets.

Praesidium/Security Service safeguards passwords with Kerberos-based security mechanisms that authenticate users and application servers through an exchange of encrypted "tickets." A ticket is a complex encrypted identification that substitutes for a password. Short-lived tickets—not valuable passwords—are transmitted throughout the network.

Security Service supports HP-UX Internet services, which form the basis of Internet/intranet functions such as ftp, remote shell (remsh), remote copy (rcp), remote login (rlogin), and telnet. Applications based on Kerberos, such as those found in banking and telecommunications industries, also can use Security Service for authentication.

Security Service provides scalable and highly available authentication and security-enhanced communication. It can be monitored through the graphical OpenView tool HP DE/Service Monitor.

HP's VirtualVault addresses the weak

link in security by providing a security gateway to the enterprise and granting selective access to sensitive information. At the foundation of VirtualVault is a trusted operating system that exceeds strict U.S. government B1 and European E3 standards and has been used in military agencies' computer networks. VirtualVault provides data partitioning and least privilege (no root) to ensure a more secure environment for Web applications.

All electronic transactions bound for sensitive internal applications must travel through VirtualVault's trusted gateway, which acts as a safeguard against users and applications gaining unauthorized access to information or systems and provides support for VirtualVault administration. Administration of the Web server and the operating system is simplified through a browser-based interface. VirtualVault also incorporates a Netscape Web server to handle external Web requests.

Firewalls provide necessary protection between organizations and the Internet, filtering common network services such as e-mail, telnet, and ftp services. They do not provide protection for the Web transaction itself. Firewalls complement VirtualVault and are used side-by-side in an architecture that connects to the Internet.

Praesidium/Security Service is priced from \$3,100 to \$15,500 per server. VirtualVault is available worldwide immediately. A single copy costs \$75,000; two to five VirtualVault licenses cost \$65,000 each; and six or more copies cost \$45,000 each.

PE/WorkManager for Workgroups 3.0

HP has announced HP PE/WorkManager for Workgroups 3.0, a new release of its team-oriented data-management solution for engineering workgroups. The new release adds

features for managing change orders and bills of materials and for classifying objects used in engineering designs.

The workgroup solution is based on the company's product-data and workflow management system, HP WorkManager, and integrates with the HP PE/ME10 2-D design and drafting application and with the HP PE/SolidDesigner solids-modeling software.

The new change-management feature helps save time and money by simplifying the setup of change orders and requests and managing the flow of information. The classification interface, introduced with this version, reduces design time by letting engineers classify, find, and retrieve objects from the database and load them directly into HP PE/ME10 or HP PE/SolidDesigner. The enhanced bills-of-materials management feature—available as an add-on module—enables multilevel scanning of the parts structure and allows updating of existing product structures directly from HP PE/ME10 and HP PE/SolidDesigner solutions.

HP PE/Workmanager for Workgroups supports various relational databases, such as Oracle, HP ALLBASE, Sybase, and Informix. It can be extended with HP WorkManager modules for additional functionality. An integration toolkit and certification program is available to partners interested in integrating their applications into the HP PE/WorkManager for Workgroups interface.

HP PE/WorkManager for Workgroups 3.0 is available now for HP and Silicon Graphics UNIX workstations, as well as for PCs running Windows. Prices start from \$1,950 per seat for HP PE/WorkManager for Workgroups, including integration to HP PE/ME10

or HP PE/SolidDesigner, and prices start from \$1,390 per seat for the integration toolkit.

Netscape SuiteSpot

HP has announced that it has signed an agreement with Netscape to include Netscape SuiteSpot server software on the HP 9000 enterprise Web server product lines. HP also announced it is shipping the following—Netscape Enterprise Server, Netscape Proxy Server, and Netscape FastTrack Server software—on its HP 9000 enterprise Web server product line.

Netscape SuiteSpot is an integrated suite of Web servers for building corporate intranets. It includes Netscape's LiveWire Pro visual development tool, plus any combination of the five Netscape SuiteSpot servers: Netscape Enterprise Server, Netscape Proxy Server, Netscape News Server, Netscape Mail Server, or Netscape Catalog Server.

Netscape uses the HP 9000 Model K400 four-way Enterprise Web Server to service its own home page, <http://www.netscape.com>. In a five-day test performed by HP and Netscape in July, the HP 9000 Enterprise Web Server handled more than 15 million hits a day, making it the highest-performing single system of any of the systems currently running on Netscape's Web site.

HP also offers consulting expertise for Netscape products within the Internet Services Practice of HP's Professional Services Organization.

Netscape FastTrack Server is available on HP 9000 Enterprise Web Servers for \$295. Netscape Enterprise Server and Netscape Proxy Server are available now on HP 9000 Enterprise Web Servers for \$995 each. Existing customers may upgrade immediately. HP plans to ship the balance of Netscape SuiteSpot—

Netscape News Server, Netscape Mail Server, and Netscape Catalog Server—beginning in the fourth quarter of 1996.

Software Update Tools

HP has announced Software Update Manager, an Internet-based online tool that makes it easier for users to support and update their software applications. Software Update Manager is a free-of-charge service available to users with support contracts for the HP 9000 Series 700 and 800.

Previously, users automatically received regular updates via postal mail and then decided whether or not to use the updates. With the Software Update Manager, they are automatically notified of available updates but receive only the updates that they order online. Software Update Manager also may be used to check the shipping status of requested updates, view comprehensive update histories on individual systems, and review data on available application and software updates.

HP 9000 Series 700 and 800 computer users may use the program as part of their current HP software support agreement at no extra charge. All users need is Internet access with Web Browser capabilities and CD-ROM as their media. Users may access the program at <http://ussupport.external.hp.com> or via HP's home page. The complete capabilities of Software Update Manager are available now in the U.S. and Canada. HP plans to make the program available in Europe by fall of 1996 and in Asia in 1997.

Attention vendors: New product announcements should be sent to New Products Editor, hp-ux/usr magazine, Interex, P.O. Box 3439, Sunnyvale, California 94088-3439, USA, or e-mail: pollace@interex.org.

Deadline for submission is two months prior to publication.



interex

Shared Knowledge.
Shared Power.

The International
Association of
Hewlett-Packard
Computing
Professionals

Membership Levels

Membership Levels

■ associate level

includes the following services/benefits:

- Subscription to either:
hp•ux/usr magazine –
includes companion Product Directory
Interact magazine –
includes companion Product Directory
- Member rates for Interex Conferences
- Read-only access to Interex Online library
- Membership in your local Regional User Group (RUG) at RUG membership rate

■ contributing level

includes the following services/benefits:

- Subscription to BOTH:
hp•ux/usr magazine –
includes companion Product Directory
Interact magazine –
includes companion Product Directory
- Subscription to InterexPress, monthly news publication
- E-Mail account through Interex (includes read-only access to Interex Online library)
- Access to Special Interest Groups (SIGs)
- Member rates for Interex Conferences
- Membership in your Regional User Group (RUG) at RUG membership rate
- Voting Privileges for Board Elections and Advocacy Surveys (i.e., system improvement surveys)

Service Package

■ online service package

includes ALL the benefits of Contributing Level plus:

- **Software Access**—unlimited downloads from entire library of HP-UX, MPE, and RTE programs (containing over 4,800 programs). Includes one free tape of current Interex Annual Release in the operating system of your choice. Custom tapes from software library are also available.

- **Information Access**—full text search and downloading capabilities for all Interex publications including: hp•ux/usr, Interact, Vendor Resource Directories, product news and announcements, and Conference Proceedings abstracts.

- **Member Access**—member directory. Find members with similar interests. Plus access to the *Who's Who* guide of Interex staff, volunteers, and HP liaisons.

Online services are continually upgraded and modified; services are subject to change without notice.

I'd Like to Join Interex

Choose one of the following

■ membership levels & service package

- ☐ Associate Level, \$49.50*
- ☐ Contributing Level, \$115.00*
- ☐ Contributing Level Plus Online Package \$595.00*

All membership and service packages are based on an annual fee.

Package subscribers, please choose the following:

■ preferred software format and operating system for the annual tape release

- ☐ 1600 BPI Magnetic Tape
- ☐ 6250 Magnetic Tape
- ☐ Linus Cartridge Tape (CS-80)
- ☐ DAT 4mm

(Check one of the following)

- ☐ HP-UX ☐ MPE/iX
- ☐ MPE V ☐ RTE

■ member directory

Please include me in the member directory.

- ☐ Yes ☐ No

Please initial: _____

■ service agreement

If you relocate, should services transfer with you?

- ☐ Yes ☐ No

Please initial: _____

■ mailing list

Would you like to receive mailings about other computer-related vendors' products and services?

- ☐ Yes ☐ No

Order Form

■ mailing address (attach business card here)

name _____

job title _____

company _____

address _____

city/state/zip/country _____

telephone/extension _____

fax _____

e-mail _____

■ payment options

- ☐ Bill me ☐ Check enclosed, payable to Interex
- ☐ Purchase order enclosed, PO# _____
(purchase order accepted for invoicing purposes only)
- ☐ Please charge my: ☐ Visa ☐ MasterCard ☐ AmEx

credit card number / expiration date _____

signature _____

Total payment enclosed \$ _____

Foreign currency accepted BUT payment must be equivalent of U.S. currency. Each publication has an annual subscription value of \$49.50.

NOTE: Services do not begin until payment is received.

*Canada & Mexico add \$25 and all other countries outside the U.S. add \$50 for additional postage.

■ send form and payment to:

Interex, P.O. Box 3439,
Sunnyvale, CA 94088-3439, USA;

fax: 408 747-0947

phone: 800.INTEREX, 408.747.0227

e-mail: membership@interex.org

CompuServe: 76376,1222

http://www.interex.org

Please read and sign the following disclaimer: I am applying for services with Interex. I understand that no funds will be returned after any Contributed Software Library (CSL) tapes/disks have been shipped or downloaded from Interex. I agree not to distribute software to any unauthorized users or use software received through Interex on more than one system at a time. I understand that this agreement stays in force even after my services expire or are terminated.

Contact the Membership Department for pricing of the Right-to-Copy License for multiple machine usage

Signature _____

Date ____/____/____

Sign up NOW for Spring 1997 Listings

hp-ux/resource directory

The *hp-ux/resource directory* is a complete resource guide for HP-UX users seeking answers. This is one of the industry's most extensive reference guides for HP-UX products, services, and vendors. It will be devoted entirely to HP 9000 users operating in multi-user, workstation, and multi-system UNIX environments. This bi-annual directory, published each year in March and September, is a separate publication mailed out with *hp-ux/usr* magazine, the only HP-specific publication on the market.

Added BONUS: your message will reach your customers for one full year on the *Internet*. Look for the directory on the Interex home page <http://www.interex.org>. The investment for a full year listing in the *hp-ux/resource directory* is \$475.

PRODUCT CATEGORIES

Accounting
Accounting Software
Alphanumeric Paging Software
Application Development Software
Application Development Tools
Application Development Tools/4GL
Application Engineering
Backup/Restore
Backup Software
Bar Code Data Collection Systems
Batch Job Management
Books
Business-Critical Application Development & Deployment
Business Software
CD-R
CAD Software/Hardware
Change Management for Software Development
Change Management Tools
Checkpoint Restart Facility
Client-Server
Client-Server Software
Communications
Communications Servers
Communications Software
Consulting
Consulting/Systems Integration
Customer Support
Customer Support/Help Desk Systems
Database Management Systems
Database Management Tools
Data Center Management
Data Migration Tool
Data Warehousing
Decision Support Systems
Diagramming & Flowcharting
Disaster Recovery
Distributed Computing
Distribution Software
Distributor
Document Management
Electronic Data Interchange (EDI)
Electronic Form Printing
E-Mail & Directory Integration
End-User Access Tools
End-User Computing
Equipment
Executive Information Systems
Facility Maintenance Software
Fax Automation
Fax Software
File Manager Utility
Financial

Forestry
Fourth Generation Language
GIS (Geographic Information System)
Government & Utility Software
Graphics
Groupware
Hardware
Hardware/Mass Storage
Hardware Subsystems
Help Desk Management
Human Resources & Personnel Systems
Image Processing
Image Storage & Retrieval Management
Industrial Terminals
Input Devices
Instrument Control
Integration Tools
Internet
Internet Commerce
Internet/Intranet
Internet Services
Internet Solutions
Inventory Control
I/O Boards
Job Scheduling & Workload Management
Justice Software
Laser Printing Software
Maintenance
Manufacturing Software
Mass Storage
Mass Storage Peripherals
Math Library
Memory
Memory Upgrades
Middleware
Migration Services
Migration Services/Tools
Migration Tools
Modular Mass Storage
Multimedia
Network Backup Software
Networking
Networking Systems
Network Management
Output Management
Payroll
PC Card Reader
PC Compatibility
PC Integration
Performance
Performance Software
Personal Information Manager
Personnel Management
Pointing Devices

Power Protection & Conditioning
Print Management
Print Management Software
Process Control Software
Production Planning
Project Management
Programmer's Editor
Protocol Converters/Interfaces-Hardware
Publications
Public Safety Software
Quality Assurance Tools
Records Management
Rentals
Report Viewing, Printing, & Distribution
Report Writers
Sales & Marketing
Scheduling
Scheduling/Task Management
Security
Service Repairs
Software
Software Backup
Software Development Tools
Software Distribution Tools
Software Maintenance & Testing
Spoolers
Spreadsheets
Statistics/Data Analysis
System Integration
System Management
System Management Tools
System Printers
3-D Graphics Tool Kit
3-D Porting Tool
Tape Backup Products
Tape Storage/Data Interchange
Technical Documentation/Cross-Referencing
Terminals
Terminal Emulation
Text & Information Retrieval
Text Editors
Time & Billing
Time Reporting Terminals
Training
User Groups
Video/Keyboard/Mouse Extension
Warehouse & Distribution Management
Workstations

Other categories may be created as needed.



hp-ux/resource directory

Spring 1997 Listing Form

	FIRST LISTING	EACH LISTING THEREAFTER	TOTAL
<input type="checkbox"/> Listing (includes two issues)	\$475	\$375	\$ _____
<input type="checkbox"/> Hyperlink to your home page	\$500	\$150	\$ _____
<input type="checkbox"/> \$1.00 Per Word Over 75 Words			\$ _____
<input type="checkbox"/> Company Product Logo or Photo	\$100	\$ 50	\$ _____
<input type="checkbox"/> Cross Reference	\$200	\$200	\$ _____

Total \$ _____

Closing Date: Thursday, December 12th, 1996

1.	_____	_____	_____
	Category	Product Name	Operating Environment
2.	_____	_____	_____
	Category	Product Name	Operating Environment
3.	_____	_____	_____
	Category	Product Name	Operating Environment
4.	_____	_____	_____
	Category	Product Name	Operating Environment
5.	_____	_____	_____
	Category	Product Name	Operating Environment

Product Description (MAXIMUM 75 WORDS) There is a \$1.00 per word charge for all listings over 75 words.

Company Name _____

Address _____

City _____ State _____

Zip/Postal Code _____ Country _____

Telephone _____ Web URL _____ Fax _____

Authorization:

Signature _____ Print Name _____

Title _____ Date _____

Please mail or fax completed form to: Interex, 1192 Borregas Avenue, Sunnyvale, CA 94088, U.S.A.,
Attention: Kathie Schwartz, 800.468.3739, ext. 620, 408.747.0227, **Fax 408.747.0947**

PAYMENT OR PURCHASE ORDER MUST ACCOMPANY ALL ORDERS

Terms & Condition of Sale: All ads are under the terms Net 30 Days. Any balance over 30 days old will be charged 1.5% per month.

Cancellations Clause: All cancellations must be made by the listing due date and must be followed in writing within 5 working days.

Artwork requirements available upon request.

Advertiser's Index

Please call or fill out adjacent card for further product information.


READER SERVICE NUMBER	ADVERTISER	PAGE #
4	Bering Technologies, Inc.	51
	(800) 237-4641 or (408) 363-6500	
6	CLAM Associates	47
	(617) 621-2542 or e-mail: marketing@clam.com	
110,111	Computer Solutions	37,77
	(407) 649-0123/Fax (407) 649-1407	
28	Concorde Technologies	7
	(800) 359-0282/Fax: (619) 536-5500	
37	Confluent	79
	(800) 780-2838 or email: info@confluent.com	
119	Consan	C-4
	(612) 949-0053/Fax (612) 949-0453	
140	C.S.U. Industries, Inc.	15
	(516) 239-4310/Fax: (516) 239-8374	
76	Dallastone	19
	(603) 647-8168/Fax: (603) 624-2466	
44	Design 3000 Plus, Inc.	25
	(503) 585-0512/ Fax: (503) 585-1706	
20	Digital StorageWorks	9
	(800) 663-1966 or (603) 890-6700	
8	Herstal Automation	13
	(941) 358-2001/Fax: (941) 358-2010	
17	IEM, Inc.	75
	(970) 221-3005 or (800) 321-4671/e-mail: info@iem.com	
171	I/O Data Systems, Inc.	65
	(216) 835-2211/Fax: (216) 835-0220	
117	ITAC Systems, Inc. - Mouse-Trak	5
	(214) 494-3073/Fax: (214) 494-4159	
129	Lund Performance Solutions	53
	(541) 926-3800	
7	Monterey Bay Communications	1
	(408) 429-6144/Fax: (408) 429-1918	

READER SERVICE NUMBER	ADVERTISER	PAGE #
71	Norco Computer Systems, Inc.	21
	(800) 892-1920/Fax: (216) 572-0636	
72	ORBIT Software	41
	(800) 6-Online or (510) 837-4143	
97	Pericom, Inc.	2
	(609) 588-5300/Fax: (609) 588-8906	
91	Personal Productivity Tools	53
	(708) 620-5000/Fax: (708) 691-0718	
100	Robelle Consulting Ltd.	57
	(604) 582-1700/Fax: (604) 582-1799	
138	SAIC	C-2
	(800) 447-4373/Fax: (619) 552-5253	
126	Sejus Corporation	C-3
	(503) 638-9000/Fax: (503) 638-9009	
180	Software Licensing Corp.	19
	(702) 832-0881	
103	Straightline	34,35
	(206) 865-8314	
175	Syncsort	17
	(201) 930-8200/Fax: (201) 930-8290	
123	Taurus Software	3
	(415) 961-1323, ext. 100	
77	Technical & Scientific Application	33
	(800) 422-4872/Fax: (713) 935-1555	
49	Ted Dasher & Associates	29
	(800) 638-4833/Fax: (205) 591-1108	
40	VESOFT, Inc.	29
	(310) 282-0420/ Fax: (310) 785-9566	
128	World Data Products	69
	(800) 553-0592/Fax: (612) 476-1903	



The New Vikon.

Full strength UNIX on-the-go!



- ☐ 13" 1280 x 1024 VUEpanel™
- ☐ 100 MHz HP PA-RISC
- ☐ 192 Mb RAM, 10 Gb disk
- ☐ Three internal bays
- ☐ DC power & rackmount option


Have you been yearning for an easy to tote system with *full* UNIX capabilities? The new Vikon™ is your answer. It's a Hewlett-Packard PA-RISC powered SuperTransportable® with a *real* workstation display.

It's for savvy, on-the-go professionals who need a TransportableWorkstation™ for on-the-site solutions, or those who simply need *full* strength™ UNIX and portability. The Vikon is small enough and light enough to go wherever you need to be. And, when you get there, you'll have a *real* UNIX workstation at your command.

The VUEpanel display will reward you with crisp, brilliant images. PA-RISC delivers power and performance, while HP-UX provides the application solution base you're looking for. And, it's portable!

So, when you're on-the-go and need a full strength UNIX platform by your side, the Vikon is your answer. It's from SEJUS.® A Hewlett-Packard authorized VAR, serving the HP community since 1986. Call us now for more information.

SEJUS... providing better solutions.®




SEJUS

SEJUS CORPORATION
2618 PALISADES CREST DRIVE
LAKE OSWEGO, OR 97034-7550 USA
PHONE 503.638.9000 FAX 503.638.9009

CIRCLE 126 ON READER SERVICE CARD

® registered, ™ trademark of SEJUS Corporation. All other trademarks are the property of their respective owners.
Printed in the USA. © 1995 SEJUS CORPORATION. All rights reserved. 4703R1295.

SATISFY your everEXPANDING needs for Data STORAGE



THINK BIG. **THINK MAMMOTH.**

THE MAMMOTH 8MM CARTRIDGE TAPE SUBSYSTEM FROM EXABYTE USES HIGH—PERFORMANCE IDRC DATA COMPRESSION to give you double capacity and throughput. That high capacity along with ultra-fast transfer rate and search speed offers the performance you require to meet growing data storage needs.

BIG ADVANTAGES—

- **HIGH PERFORMANCE**—for digital video and multimedia applications
- **HIGH CAPACITY**—40GB [compressed] at 6MB/sec. transfer rate
- **COMPATIBILITY** with earlier 8mm releases.

The Mammoth 8MM tape subsystem is available today from Consan. Consan distributes a full line of tape, disk and optical storage products from the world's leading manufacturers. For more information on the Mammoth subsystem or other storage products, call Consan today at **1-800-229-DISK**

[3 4 7 5].




EXABYTE®
We're Backing It Up.

CONSAN

7676 EXECUTIVE DRIVE
EDEN PRAIRIE, MN 55344
TEL: 612-949-0053
FAX: 612-949-0453

1320 TOWER ROAD
SCHAUMBURG, IL 60173
TEL: 847-519-1060
FAX: 847-519-1248

101 EAST PARK BLVD.
PLANO, TX 75074
TEL: 214-422-3392
FAX: 214-422-3397



© 1996 CONSAN, INC.

CIRCLE 119 ON READER SERVICE CARD